



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁵ :

H04B 7/216, H04L 12/46

A1

(11) International Publication Number:

WO 94/27382

(43) International Publication Date: 24 November 1994 (24.11.94)

(21) International Application Number: PCT/US94/05037

(22) International Filing Date: 6 May 1994 (06.05.94)

(30) Priority Data:

08/062,457	11 May 1993 (11.05.93)	US
08/076,340	11 June 1993 (11.06.93)	US
08/085,662	29 June 1993 (29.06.93)	US
08/101,254	3 August 1993 (03.08.93)	US
08/147,377	3 November 1993 (03.11.93)	US
08/168,478	16 December 1993 (16.12.93)	US
PCT/US93/12628	23 December 1993 (23.12.93)	WO
(34) Countries for which the regional or international application was filed: AT et al.		
08/198,452	18 February 1994 (18.02.94)	US
08/198,404	22 February 1994 (22.02.94)	US
08/205,689	4 March 1994 (04.03.94)	US

(74) Agent: WINSLADE, Christopher, C.; McAndrews, Held & Malloy Ltd., Suite 3400, 500 West Madison Street, Chicago, IL 60661 (US).

(81) Designated States: AU, CA, FI, NO, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).

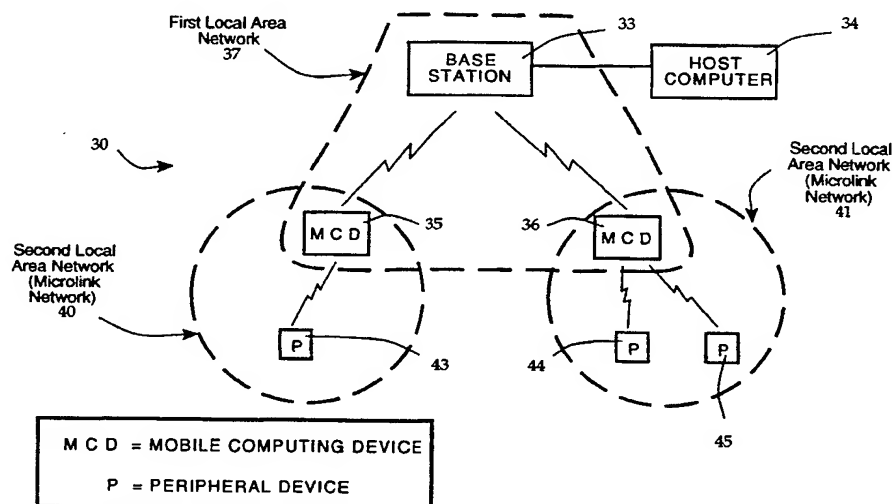
Published

*With international search report.**Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.*

(71) Applicant: NORAND CORPORATION [US/US]; 550 Second Street Southeast, Cedar Rapids, IA 52401 (US).

(72) Inventors: MAHANY, Ronald, L.; 1330 Sierra Drive NE, Cedar Rapids, IA 52402 (US). WEST, Guy, J.; 1829 Grande Avenue SE, Cedar Rapids, IA 52403 (US). BUNTE, Alan, G.; 5510 Hickory Wood Court, Cedar Rapids, IA 52402 (US).

(54) Title: PREMISES, PERIPHERAL AND VEHICULAR LOCAL AREA NETWORKING



(57) Abstract

A communication system (30) in which multiple local area networks (37, 40, 41) exhibiting different characteristics are employed to link portable or mobile computing devices (35, 36). Radio base stations (33, 56, 57, 58, 59) connected to a backbone LAN and at least one mobile computing device (35, 36) form a high-power LAN (37), using a frequency hopping protocol. A low-power LAN (40, 41) allows for radio communication between a mobile computing device (35, 36) and peripheral devices (43, 44, 45), using a single-frequency spread spectrum protocol. A vehicular LAN provides for short-range communication between vehicle terminals (3007) and portable terminals (3009). Radio units (35, 36) participating in multiple LANs (37, 40, 41) use the appropriate protocol, frequency and power level for communicating through the communication system (30).

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australia	GE	Georgia	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KE	Kenya	RO	Romania
CA	Canada	KG	Kyrgystan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LI	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LU	Luxembourg	TD	Chad
CS	Czechoslovakia	LV	Latvia	TG	Togo
CZ	Czech Republic	MC	Monaco	TJ	Tajikistan
DE	Germany	MD	Republic of Moldova	TT	Trinidad and Tobago
DK	Denmark	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	US	United States of America
FI	Finland	MN	Mongolia	UZ	Uzbekistan
FR	France			VN	Viet Nam
GA	Gabon				

PREMISES, PERIPHERAL AND VEHICULAR LOCAL AREA NETWORKING

BACKGROUND OF THE INVENTION

This application is a continuation-in-part of U.S. Application Ser. No. 08/198,404, filed February 22, 1994 by Mahany et al. (Attorney Docket Nos. 10082US07; DN37998XA) which is a continuation of U.S. Application Ser. No. 08/198,452, filed February 18, 1994 by Mahany et al. (Attorney Docket Nos. 10082US06; DN37998X), which is in turn a continuation-in-part of U.S. Application Ser. No. 08/168,478, filed December 16, 1993 by R. Mahany (Attorney Docket Nos. 10092US06 and DN37998E), which is itself a continuation-in-part of U.S. Application Ser. No. 08/147,377 filed November 3, 1993 by R. Mahany (Attorney Docket No. DN37998D), which is in turn a continuation-in-part of U.S. Application Ser. No. 08/101,254 filed August 3, 1993 by R. Mahany (Attorney Docket No. DN37998C), which is in turn a continuation-in-part of U.S. Application Ser. No. 08/085,662 filed June 29, 1993 by R. Mahany (Attorney Docket No. DN37998B), which is in turn a continuation-in-part of U.S. Application Ser. No. 08/076,340 filed June 11, 1993 by R. Mahany (Attorney Docket No. DN37998A), which is in turn a continuation-in-part of U.S. Application Ser. No.

- 2 -

08/062,457 filed May 11, 1993 by R. Mahany (Attorney Docket No. DN37998).

U.S. Application Ser. No. 08/198,452 filed February 18, 1994 by Mahany et al. (Attorney Docket Nos. 10082US06; DN37998X), is also a continuation-in-part of PCT Application Ser. No. PCT/US93/12628 filed December 23, 1993 by Mahany et al. (Attorney Docket Nos. DN37967C and 10082WO01), which is itself a continuation-in-part of pending U.S. Application Ser. No. 08/027,140 filed March 5 1993 by Mahany et al. (Attorney Docket Nos. DN37967B and 10082US05), which is itself a continuation-in-part of U.S. Application Ser. 07/997,693 filed December 23, 1992 by Luse et al. (Attorney Docket Nos. DN37967A and 10005US02), now abandoned, which is itself a continuation-in-part of U.S. Application Ser No. 07/982,292 filed November 27, 1992 by Luse et al. (Attorney Docket Nos. DN37967 and 92 P 837), now abandoned, which is itself a continuation-in-part of U.S. Application Ser. No. 07/700,704 filed May 14, 1991 by Mahany et al. (Attorney Docket Nos. DN37834X and 91P383), now abandoned, which is itself a continuation-in-part of U.S. Application Ser. No. 07/699,818 filed May 13, 1991 by Mahany et al. (Attorney Docket Nos. DN37834 and 91P862), now abandoned.

This application is also a continuation-in-part of U.S. Application Ser. No. _____ filed March 4, 1994 by Mahany et al. (Attorney Docket Nos. DN37139XXA and 10458US03) which is a continuation-in-part of U.S. Application Ser. No. 07/735,128 filed July 22, 1991 by Koenck et al. (Attorney Docket Nos. DN37139XX and 91P326) which is itself a continuation-in-part of U.S. Application Ser. No. 07/467,096 filed January 18, 1990 by Koenck et al. (Attorney Docket Nos. DN37139), now U.S. Patent No. 5,052,020.

- 3 -

INCORPORATION BY REFERENCE

The above referenced applications, PCT Application No. PCT/US92/08610 filed October 1, 1992, as published under International Publication No. WO 93/07691 on April 15, 1993, together with U.S. Patent No. 5,070,536, by Mahany et al., U.S. Patent No. 4,924,426, by Sojka, and U.S. Patent No. 4,910,794, by Mahany, are incorporated herein by reference in their entirety, including drawings and appendices, and hereby are made a part of this application.

- 4 -

TECHNICAL FIELD

The present invention relates generally to local area networks used for transmitting and receiving information and more particularly to a singular radio using multiple communication protocols for servicing corresponding multiple radio local area networks.

BACKGROUND OF THE INVENTION

Multiple radio base station networks have been developed to overcome a variety of problems with single radio base station networks such as spanning physical radio wave penetration barriers, wasted transmission power by portable computing devices, etc. However, multiple radio base station networks have their own inherent problems. For example, in a multiple base station network employing a single shared channel, each base station transmission is prone to collision with neighboring base station transmissions in the overlapping coverage areas between the base stations. Therefore, it often proves undesirable for each base station to use a single or common communication channel.

In contradistinction, to facilitate the roaming of portable or mobile devices from one coverage area to another, use of a common communication channel for all of the base stations is convenient. A roaming device may easily move between coverage areas without loss of connectivity to the network.

Such exemplary competing commonality factors have resulted in tradeoff decisions in network design. These factors become even more significant when implementing a frequency hopping spread spectrum network. Frequency hopping is a desirable transmission technique because of its ability to combat frequency selective fading, avoid narrowband interference, and provide multiple communications channels.

- 5 -

Again, however, changing operating parameters between coverage areas creates difficulties for the roaming devices which move therebetween. In particular, when different communication parameters are used, a portable or mobile device roaming into a new base station coverage area is not able to communicate with the new base station without obtaining and synchronizing to the new parameters. This causes communication backlog in the network.

Moreover, even when a radio frequency network is established to cover the premises of a building or group of buildings, certain types of communication flow between certain types of devices make for inefficient use of such a network. In fact, an ordinarily efficient network configuration may be deemed intolerable in certain communication scenarios.

Computer terminals and peripheral devices are widely used. Many types of computer terminals exist which vary greatly in terms of function, power and speed. Many different types of peripheral devices also exist, such as printers, modems, graphics scanners, text scanners, code readers, magnetic card readers, external monitors, voice command interfaces, external storage devices, and so on.

Computer terminals have become dramatically smaller and more portable, as, for example, lap top computers and notebook computers. Computer terminals exist which are small enough to be mounted in a vehicle such as a delivery truck or on a fork lift. Hand held computer terminals exist which a user can carry in one hand and operate with the other.

Typical computer terminals must physically interface with peripheral devices. Thus, there must either be a cable running from the computer terminal to each peripheral device, or the computer terminal must be

- 6 -

docked with the device while information transfer takes place.

In an office or work place setting, the physical connection is typically done with cables. These cables
5 pose several problems. For example, many cables are required in order for a computer terminal to accommodate many peripheral devices. In addition, placement of peripheral devices is limited by cable lengths. While
longer cables may be used, they are costly.
10 Additionally, there may be a limited number of ports on a computer terminal, thus limiting the number of peripherals that may be attached.

Another problem arises when several computer terminals must share the same peripheral device, such as
15 a printer. All of the computers must be hardwired to the printer, which may create a protocol problem if the computer terminals are of different types.

Peripheral cabling is an even greater problem in scenarios where hand-held and portable computer terminals
20 are used. The cabling required for an operator to carry a hand-held computer terminal in one hand, have a small portable printer attached to his belt, and carry a code reader in the other hand is cumbersome and potentially even dangerous. For example, such an operator loses a
25 great deal of mobility and flexibility while supporting a number of cabled devices. In addition, as cables wear out and break, exposed electric current could shock the operator, or create a spark and potentially cause a fire or explosion in some work areas.

30 The requirement of physically connecting the computer terminals and peripherals severely reduces the efficiency gained by making the devices smaller. An operator must somehow account for all of the devices in a system and keep them all connected. This can be very
35 inconvenient. For example, an operator having a notebook

- 7 -

computer and a modem in a briefcase may wish to have the freedom to move around with the computer but without the modem. He may, for example, wish to work at various locations on a job sight and at various times transmit or receive information via his modem. If the modem and the computer are hard wired together, he must either carry the modem with him or keep connecting and disconnecting it.

Furthermore, cabling can be expensive because cables frequently prove to be unreliable and must be replaced frequently. In portable environments, cables are subject to frequent handling, temperature extremes, dropping and other physical trauma. It is not uncommon for the cables or the connectors for the cables on the devices to need replacing every three months or so.

Attempts to alleviate or eliminate these problems have been made but have not been entirely successful. One solution is to incorporate a computer terminal and all of the peripherals into one unit. However, this solution proves unsatisfactory for several reasons. For example, the incorporation of many devices into one unit greatly increases the size and weight of the unit, thus jeopardizing its portability. Furthermore, incorporating all of the functions into one unit greatly reduces and, in most cases eliminates, the flexibility of the overall system. A user may only wish to use a hand-held computer terminal at times, but at other times may also need to use a printer or occasionally a code reader. An all-incorporated unit thus becomes either overly large because it must include everything, or very limiting because it does not include everything.

Another solution has been to set up Local Area Networks (LAN's) utilizing various forms of RF (Radio Frequency) communication. The LAN's to date, however, have been designed for large scale wireless

- 8 -

communications between several portable computer terminals and a host computer. Therein, the host computer, itself generally a stationary device, manages a series of stationary peripherals that, upon requests to the host, may be utilized by the portable terminals. Other large scale wireless communications have also been developed which provide for RF communication between several computer terminals and peripheral devices, but have proven to be ineffective as an overall solution. For example, these systems require the peripheral devices to remain active at all times to listen for an occasional communication. Although this requirement may be acceptable for stationary peripheral devices receiving virtually unlimited power (i.e., when plugged into an AC outlet), it proves detrimental to portable peripherals by unnecessarily draining battery power. Similarly, in such systems, the computer terminals are also required to remain active to receive an occasional communication not only from the other terminals or the host, but also from the peripherals. Again, often unnecessarily, battery power is wasted.

In addition, such large scale systems are designed for long range RF communication and often require either a licensed frequency or must be operated using spread spectrum technology. Radios in such systems are typically cost prohibitive, prove too large for convenient use with personal computers and small peripheral devices, and require a great deal of transmission energy utilization.

Furthermore, these systems do not provide for efficient communication between portable computer devices and peripherals. For example, a portable computer device may be mounted in a delivery truck and a driver may desire to transmit data to, or receive data from, a host computer or peripheral device at a remote warehouse

- 9 -

location. While permitting such transmission, such wide area networks (WANs) only provide point-to-point communications, use a narrow bandwidth, and often have heavy communication traffic. As a result, WANs are generally slow and expensive and simply do not provide an effective overall solution.

Additionally, in order for a computer device to be effectively portable in these systems, it must be capable of participating on any number of LANs operating with different communication parameters and protocols. Thus, each portable computer device requires a plurality of built-in radio transceivers, one to accommodate each of such LANs. As a result, portable computer devices can become costly, excessively large, heavy, and power hungry.

Thus, there is a need for a radio frequency communication system and associated radio that supports the use of network peripherals and solves the foregoing problems relating to power conservation and portability.

Another object of the invention is to provide a method and apparatus wherein collisions are minimized in overlapping coverage areas in a multiple base station network while providing a seamless communication network to support roaming devices.

Yet another object of the invention is to provide a method and apparatus wherein collisions are minimized in overlapping coverage areas by utilizing uncommon communication channel characteristics in a multiple base station network, while still providing seamless communication for roaming devices by informing roaming devices of the nature of the neighboring base station communication channel characteristics.

A further object of the invention is to provide a communication network wherein base stations communicate

- 10 -

with roaming devices to help conserve transmission power usage.

5 A still further object of the present invention is to provide a hierarchical communications system utilizing spread spectrum frequency hopping communication.

Other objects, advantages, and novel features of the present invention will become apparent from the following detailed description of the invention when considered in conjunction with the accompanying drawings.

- 11 -

SUMMARY OF THE INVENTION

The present invention solves many of the foregoing problems in a variety of embodiments. The network and associated radio of the present invention provides wireless peripheralization of roaming computing devices and data collection devices. The roaming computing devices communicate over an extended area via a first local area network. The first local area network is a high power radio communications system. Each mobile computer device can communicate with peripheral devices via a second, low-power local area network. Additionally, the roaming computing device, also called the parent device, and peripherals may communicate within a limited area while moving within an independent wireless network that provides coverage over a much broader service area. Thus, the communications system comprises at least two independent wireless networks, with the parent device participating in the multiple networks by selectively processing and controlling the flow of information among devices connected to the multiple networks.

In some embodiments, roaming computing or data collection devices and their peripherals communicate within a building or group of buildings serviced by a wireless LAN, hereinafter called a premises LAN. Devices that are constituents of the premises LAN may also contain facilities (transceivers and protocols) for communicating with their peripherals via the separate low power, short range radio LAN, hereinafter a peripheral LAN or MicroLAN. The parent device contains significant processing power, such that information received from various peripherals and other user input means, such as a keyboard attached to the parent unit, is combined to form a message that is communicated over the premises LAN. Information received from peripherals is

- 12 -

selectively communicated over the premises LAN in accordance with an application program, emulation mode, or operating system resident in the parent device.

5 Likewise, information received from the premises LAN may be processed and selectively forwarded to a peripheral. For example, a record received through the premises LAN may be combined with information obtained locally through another peripheral or user input means and processed by a local application to generate an
10 invoice or shipping document that is then sent wirelessly to a portable printer peripheral.

The parent device may also enable communication among peripheral devices within the peripheral LAN service area. Such communication may be either forwarded
15 from source to destination peripherals through the parent device, or directly exchanged (peer to peer communications). The former occurs if the peripherals are within communication range of the parent but not each other, or if other system design constraints, such as
20 power management dictate a centralized coordination function for power management or data transfer efficiency.

The present invention is also capable of operation within radio Wide Area Networks (WAN's). Vehicular based
25 data communication is currently serviced by a variety of public and common carrier radio WAN's that provide connectivity to computer resources anywhere in the world, such as RAM Mobile Data, ARDIS, MTEL, data PCS, CDPD, SMR, etc. The radio WAN's are generally bandwidth
30 limited, and users are charged for service on the basis of the amount of data transferred. The hierarchical network provides means of connecting a premises LAN with a group of remote devices via a radio WAN so as to minimize the expense and delays of such radio WAN's.

- 13 -

For example, communication between a parent device or radio terminal mounted in a vehicle, hereinafter a "vehicle terminal", and hand-held mobile terminal(s) that roam in an area local to the vehicle can be carried out using radio communication not subject to air time fees. Hereinafter, the network formed by the vehicle terminal and associated hand-held terminal(s) is referred to as a "vehicular LAN".

The intelligence of the parent device, in the previous example a vehicle terminal, is key in processing information generated locally in the vehicular LAN and passing only essential information through the radio WAN to minimize cost of use. Information that is not time critical may be selectively processed and stored for later batch downloading via the premises LAN using hard-wired modems or wireless communication at a docking station located at a depot or central office.

The autonomous operation of the vehicular LAN allows it to continue to function when it is out of range of the WAN, when the WAN is inaccessible during peak usage periods, or when economics dictate that WAN communication is unjustified.

The present invention relates generally to local area networks and, more specifically, to a communication system for maintaining connectivity between devices on networks which have different operating parameters while limiting the power drain of battery powered devices.

In addition, a roaming computing device may have a single radio unit which has a control processor, memory, and a transceiver. The transceiver is capable of participating in at least a first and second local area network which operate using a first and second communication protocol, respectively. The radio unit may participate as a slave to the first network pursuant to the first protocol and as a master to the second network

- 14 -

pursuant to the second protocol, and the control processor resolves conflicts between the first and second protocols.

5 In a further embodiment of the present invention, the control processor causes the radio unit to enter a state of low power consumption when the radio unit is not communicating on either the first or second network.

- 15 -

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1A is a diagrammatic illustration of a hierarchal communication system built in accordance with the present invention.

5 Fig. 1B is a diagrammatic illustration of another hierarchal communication system built in accordance with the present invention.

10 Fig. 1C is a diagrammatic illustration of still another hierarchal communication system built in accordance with the present invention.

Fig. 2 illustrates an embodiment of a basic access interval structure used by a hierarchical network of the present invention.

15 Figs. 3A and 3B illustrate the frequency of operation periodically changing corresponding to access interval boundaries in a frequency hopping communication protocol of the present invention.

20 Figs. 4A and 4B illustrate more than one access interval being used per hop in a frequency hopping communication protocol of the present invention.

Fig. 5A illustrates an embodiment of an access interval used by the hierarchical network of the present invention wherein a reservation phase is Idle Sense Multiple Access.

25 Fig. 5B illustrates an embodiment of an access interval used by the hierarchical network of the present invention wherein a device response follows a reservation poll.

30 Fig. 6A illustrates an embodiment of an access interval used by the hierarchical network of the present invention having multiple reservation slots for transmission of a Request For Poll signal.

Fig. 6B illustrates an embodiment of an access interval used by the hierarchical network of the present

- 16 -

invention wherein general devices contend for channel access.

5 Fig. 7A illustrates a sequence in an access interval used by the hierarchical network of the present invention for transferring data from a remote device to a control point device.

10 Fig. 7B illustrates a sequence in an access interval used by the hierarchical network of the present invention for transferring data from a control point device to a remote device.

Fig. 8 illustrates a preferred embodiment of an access interval used by the hierarchical network of the present invention.

15 Figs. 9A and 9B conceptually illustrate how multiple NETs may be employed in an idealized cellular-type installation according to the present invention.

Fig. 10 illustrates a base station coverage contour overlap for the multiple NETs Infrastructured Network of Fig. 1.

20 Fig. 11 illustrates hopping sequence reuse in a multiple NET configuration of the present invention.

Fig. 12 illustrates a hierarchical infrastructured network of the present invention wherein a wireless link connects base stations on separate hard wired LANs.

25 Fig. 13 illustrates a hierarchical infrastructured network of the present invention including a wireless base station.

30 Fig. 14 illustrates conceptually base stations communicating neighboring base station information to facilitate roaming of portable/mobile devices.

Fig. 15 illustrates a secondary access interval used in the MicroLAN or peripheral LAN in the hierarchical communication network according to the present invention.

- 17 -

Figs. 16a and 16b illustrate a flow chart showing the selection of a base station by a mobile computing device for communication exchange.

Fig. 17 is a flow chart illustrating a terminal maintaining synchronization with the network after it has gone to sleep for several access intervals.

Fig. 18 is a flow chart illustrating a terminal maintaining or achieving synchronization with the network after it has gone to sleep for several seconds.

Figs. 19a, 19b-1 and 19b-2 illustrate a flow chart showing an access interval during inbound communication.

Figs. 20a, 20b-1 and 20b-2 illustrate a flow chart showing an access interval during outbound communication.

Fig. 21 illustrates a sequence in an access interval used in the hierarchical communication network of the present invention with Time Division Multiple Access slots positioned at the end of the access interval.

Fig. 22 illustrates a sequence in an access interval used by the hierarchical network of the present invention with the Time Division Multiple Access slots positioned immediately following the SYNC.

Fig. 23 illustrates a sequence in an access interval used by the hierarchical network of the present invention with the Time Division Multiple Access slots positioned immediately following the SYNC and Reservation Poll.

Fig. 24 illustrates another sequence in an access interval used by the hierarchical network of the present invention with the Time Division Multiple Access slots positioned immediately following the SYNC.

Fig. 25 illustrates a portion of an access interval including the preamble, SYNC and Reservation Poll.

Fig. 26 illustrates the information contained in a sample SYNC message.

Fig. 27 illustrates the information contained in a sample Reservation Poll.

- 18 -

Fig. 28a illustrates a warehouse environment incorporating a communication network which maintains communication connectivity between the various network devices according to the present invention.

5 Fig. 28b illustrates other features of the present invention in the use of a vehicular LAN which is capable of detaching from the premises LAN when moving out of radio range of the premises LAN to perform a service, and reattaching to the premises LAN when moving within range
10 to automatically report on the services rendered.

Fig. 28c illustrate other features of the present invention in the use of a vehicular LAN which, when out of range of the premises LAN, is still capable gaining access to the premises LAN via radio WAN communication.

15 Fig. 29 is a diagrammatic illustration of the use of a peripheral LAN supporting roaming data collection by an operator according to the present invention.

Fig. 30 is a block diagram illustrating the functionality of RF transceivers built in accordance with
20 the present invention.

Fig. 31 is a diagrammatic illustration of an alternate embodiment of the peripheral LAN shown in Fig. 2.

Fig. 32 is a block diagram illustrating a channel access algorithm used by peripheral LAN slave devices in
25 accordance with the present invention.

Fig. 33a is a timing diagram of the protocol used according to the present invention illustrating a typical communication exchange between a peripheral LAN master device having virtually unlimited power resources and a
30 peripheral LAN slave device.

Fig. 33b is a timing diagram of the protocol used according to the present invention illustrating a typical communication exchange between a peripheral LAN master

- 19 -

device having limited power resources and a peripheral LAN slave device.

Fig. 33c is also a timing diagram of the protocol used which illustrates a scenario wherein the peripheral
5 LAN master device fails to service the peripheral LAN slave devices.

Fig. 34 is a timing diagram illustrating the peripheral LAN master device's servicing of both the
10 higher power portion of the premises LAN as well as the lower power peripheral LAN subnetwork with a single or plural radio transceivers.

Figs. 35 and 36 are block diagrams illustrating additional power saving features according to the present
15 invention wherein ranging and battery parameters are used to optimally select the appropriate data rate and power level of subsequent transmissions.

Fig. 37 illustrates an exemplary block diagram of a radio unit capable of current participation on multiple
LANs according to the present invention.

20 Fig. 38 illustrates an exemplary functional layout of the frequency generator of Fig. 37 according to one embodiment of the present invention.

Fig. 39 illustrates further detail of the receiver RF processing circuit of Fig. 37 according to one
25 embodiment of the present invention.

Fig. 40 illustrates further detail of the receiver signal processing circuit of Fig. 37 according to one
embodiment of the present invention.

30 Fig. 41 illustrates further detail of the receiver signal processing circuit of Fig. 37 according to another embodiment of the present invention.

Fig. 42 illustrates further detail of the memory unit of Fig. 37 according to one embodiment of the
present invention.

- 20 -

Fig. 43 illustrates a software flow chart describing the operation of the control processor in controlling the battery powered radio unit to participate on multiple LANs.

5 Fig. 44 is an alternate embodiment of the software flow chart wherein the control processor participates on a master LAN and, when needed, on a slave LAN.

- 21 -

DETAILED DESCRIPTION OF THE INVENTION

Fig. 1A illustrates a hierarchical communication system 10 within a building in accordance with the present invention. The illustrated hierarchical communication system 10 includes a local area network (LAN) for maintaining typical communication flow within the building premises, herein referred to as a premises LAN. The premises LAN is designed to provide efficient end-to-end routing of information among hardwired and wireless, stationary and roaming devices located within the hierarchical communication system 10.

The premises LAN consists of an infrastructure network comprising radio base stations 15 and a data base server 16 which may be part of a more extensive, wired LAN (not shown). The radio base stations 15 may communicate with each other via hard-wired links, such as Ethernet, RS232, etc., or via wireless (radio frequency) links. A plurality of roaming terminal devices, such as a roaming computing device 20, participate in the premises LAN of the hierarchical communication network 10 to exchange information with: 1) other roaming computing devices; 2) the data base server 16; 3) other devices which might be associated with data base server 16 (not shown); and 4) any other devices accessible via the premises LAN (not shown). A roaming computing device can be, for example, a hand-held computer terminal or vehicle mounted computer terminal (vehicle terminal).

In most circumstances, the premises LAN provides a rather optimal solution to the communication needs of a given network. However, in some circumstances, to serve a variety of particular communication needs, the premises LAN does not offer the optimal solution. Instead of relying on the premises LAN for such communications, when and where beneficial, alternate LANs are spontaneously created by (or with) network devices, such as the roaming

- 22 -

computing device 20, within the hierarchical communication system 10. Such spontaneously created LANs are referred to herein as spontaneous LANs. After the immediate benefits end, i.e., a task has been completed, or if the participants of the spontaneous LAN move out of range of each other, the spontaneous LAN terminates operation.

An exemplary spontaneous LAN involves the use of peripheral devices as illustrated in Fig. 1A. Although bulk data transfer destined for a peripheral device 23, such as a printer, from the roaming computing device 20 might be communicated through the premises LAN, a more direct interconnection proves less intrusive, saves power, and offers a lower cost solution. Specifically, instead of communicating through the premise LAN, the roaming computing device 20 needing to print: 1) identifies the presence of an available printer, the peripheral device 23; 2) establishes an RF link (binds) with the peripheral device 23; 3) directly begins transferring the bulk data for printing; and 4) lastly, when the roaming terminal finishes the transfer, the spontaneous LAN with the peripheral device 23 terminates. A spontaneous LAN created between the computing devices and peripheral devices is herein referred to as a peripheral LAN. Other types of spontaneous LANs, such as vehicular LANs, are also possible. Embodiments described below identify vehicular LANs and wide area radio networks (WANs) which are part of the hierarchical communication system according to the present invention.

Although a spontaneous LAN may operate completely independent of the premises LAN, it is more likely that there will be some degree of coordination between the two. For example, while participating in the peripheral LAN, the roaming computing device 20 may terminate participation in the premises LAN, and vice versa.

- 23 -

Alternately, the roaming computing device 20 may only service the peripheral LAN when specific participation on the premises LAN is not required, or vice versa. Moreover, the roaming computing device 20 may attempt to service each peripheral LAN as necessary in a balanced time-sharing fashion, placing little priority upon either LAN. Thus, based on the protocols and hardware selected, a spontaneous LAN can be configured so as to exist hierarchically above, below, at the same level, or independent of the premises LAN.

In generally, to design a given LAN configuration, only the characteristics of that LAN are considered for optimization purposes. However, in the hierarchical communication system of the present invention, the operation of other LANs must also be taken into account. For example, because of the roaming computing devices participation in both the premises and peripheral LANs, the requirements and operation of the premises LAN must be taken into consideration when defining the peripheral LAN, and vice versa. Thus, the hierarchical communication system of the present invention provides a series of tightly coupled radio LANs and WANS with radio transceiver and communication protocol designs which take into consideration such factors as cost, weight, power conservation, channel loading, response times, interference, communication flow, etc., as modified by a primary factor of multiple participation.

The peripheral LAN replaces hard-wired connection between a roaming computing device and associated peripherals. In a typical configuration, a peripheral LAN will consist of one or more peripherals slaved to a single master roaming computing device, although multiple master roaming computing devices are possible. Peripheral devices may be printers, code scanners, magnetic card readers, input stylus, etc.

- 24 -

Each of the peripheral devices 22 has a built-in radio transceiver to communicate with the roaming computing devices 20. The roaming computing devices 20 are configured with built-in radio transceivers capable of communicating on both the peripheral and premises LAN. The radio base stations 15 may be configured with radio transceivers only capable of communicating in the premises LAN. In alternate embodiments, as described below, the radio base stations 15 might instead be configured to participate on both the premises and peripheral LANs.

In particular, the peripheral LAN is intended to provide communications between two or more devices operating within near proximity, e.g., distances of a few tens of feet. The majority of constituents of the peripheral LAN are generally devices that do not require access to resources outside their immediate group, or which can suffice with indirect access through devices which participate outside their immediate peripheral LAN group. In contradistinction, the premises LAN is intended to provide communications between relatively many devices operating across great distances throughout a building.

The characteristics of the peripheral LAN permit the use of radio transceivers of lower cost, lower power consumption, and generally more simplistic operation than permitted by the premises LAN. However, the operation of the peripheral LAN is adapted for integration with the premises LAN so that a radio transceiver and protocol designed for operation on the premises LAN includes features which allow concurrent or sequentially concurrent operation on the peripheral LAN. For example, by selecting similar communication hardware characteristics and integrating protocols, communication

- 25 -

within the premises and peripheral LANs may be achieved with a single radio transceiver.

In one embodiment, radio communication through the premises LAN, i.e., among the radio base stations 15 and the roaming computing device 20, utilizes relatively higher-power spread-spectrum frequency-hopping communication with a reservation access protocol. The reservation access protocol facilitates frequency-hopping and supports adaptive data rate selection. Adaptive data rate selection is based upon the quality of communication on the premises LAN radio channel. Radio communication through the peripheral LAN utilizes a relatively lower-power single frequency communication also with a reservation access protocol. As more fully described below, the coordinated use of reservation access protocols in the peripheral and premises LANs maximize information flow while minimizing conflicts between devices participating in the two LANs.

Referring to Fig. 1B, a small hierarchal communication system 30 built in accordance with the present invention is shown. A base station 33 and two roaming or mobile computing devices 35 and 36 form a premises LAN 37. The premises LAN 37 provides for communication among the mobile computing devices 35 and 36 and a host computer 34. The mobile computing devices 35 and 36 can roam anywhere within the range of the base station 33 and still communicate with the host computer 34 via the base station 33.

Two peripheral LANs 40 and 41 allow for wireless communication between each mobile computing device 35 and 36 and its respective peripheral devices 43, 44 and 45 when the mobile computing device is not communicating on the premises LAN 37. Specifically, the peripheral LAN 40 consists of the mobile computing device 35 and the peripheral device 43, while the peripheral LAN 41

- 26 -

consists of the mobile computing device 36 and the two peripheral devices 44 and 45.

Fig. 1C illustrates another embodiment according to the present invention of a larger hierarchal communication system 50. The host computer 55 is connected to base stations 56, 57, 58 and 59. The host computer 55 and the base stations 56, 57, 58 and 59 provide the infrastructure for the premises LAN. The base stations need not be hard-wired together. For example, as illustrated in Fig. 1C, the base stations 56, 57 and 58 access each other and the host computer 55 via a hard-wired link, while the base station 59 accomplishes such access via a wireless link with the base station 58.

The base stations 56, 58 and 59 can support multiple mobile computing devices. For example, the base station 56 uses a frequency-hopping communication protocol for maintaining communication with mobile computing devices 61 and 62. Moreover, each of the mobile computing devices may roam out of range of the base station with which they have been communicating and into the range of a base station with which they will at least temporarily communicate. Together, the host computer 55 and the base stations 56, 57, 58 and 59 and mobile computing devices 61, 62, 64, 65 and 66 constitute a premises LAN.

More particularly, each base station operates with a different set of communication parameters. For example, each base station may use a different frequency hopping sequence. Additionally, different base stations may not employ a common master clock and will not be synchronized so as to have the frequency hopping sequences start at the same time.

Mobile computing devices 61, 62, 64, 65 and 66 are capable of roaming into the vicinity of any of the base stations 56, 58 and 59 and connecting thereto. For

- 27 -

example, mobile computing device 62 may roam into the coverage area of base station 58, disconnecting from base station 56 and connecting to base station 58, without losing connectivity with the premises LAN.

5 Each mobile computing device 61, 62, 64, 65 and 66 also participates with associated peripherals in a peripheral LAN. Each peripheral LAN is made up of the master device and its slave device. Similarly, as illustrated, the base station 57 is shown as a direct
10 participant in not only the premises LAN but also in the peripheral LAN. The base station 57 may either have limited or full participation in the premises LAN. For example, the base station 57 may be configured as a mobile computing device with the full RF capability of
15 transmission in both the premises and peripheral LANs. Instead, however, participation in the premises LAN may be limited to communicating through the hard-wired link, effectively dedicating the base station 57 to the task of servicing peripherals.

20 Although the use of a plurality of built-in radio transceivers could be used so as to permit simultaneous participation by a single device, factors of cost, size, power and weight make it desirable to only build-in a single radio transceiver capable of multiple
25 participation. Furthermore, even where a plurality of radio transceivers are built-in, simultaneous participation may not be possible depending upon the potential transmission interference between transceivers. In fact, full simultaneous participation may not be
30 desirable at least from a processing standpoint when one transceiver, servicing one LAN, always or usually takes precedence over the other. Justification for such precedence generally exists in a premises LAN over a peripheral LAN.

- 28 -

For example, communication flow in most premises LANs must be fast, efficient and rather robust when considering the multitude of participants that operate thereon. In the peripheral LAN, however, response times and other transmission related delays are generally more acceptable -- even adding extra seconds to a peripheral printer's print time will usually not bother the user. Thus, in such communication environments, it may be desirable to design the transmitters and associated protocols so that the premises LAN takes precedence over the peripheral LAN. This may yield a communication system where fully simultaneous participation in both the premises and peripheral LANs does not exist.

In communication environments wherein fully simultaneous participation does not exist or is not desired, transmitter circuitry might be shared for participation in both the premises and peripheral LANs. Similarly, in such environments, the communication protocol for the peripheral LAN can be tightly coupled with the protocol for the premises LAN, i.e., integrated protocols, so as to accommodate multiple participation. Moreover, one protocol might be designed to take precedence over the other. For example, the premises LAN protocol might be designed so as to minimize participation or response time in the peripheral LAN. As described in more detail below, such transceiver and protocol analysis also takes place when considering additional multiple participation in the vehicular LAN and WAN environments.

Fig. 2 illustrates an embodiment of a communication protocol for the premises LAN which uses a basic Access Interval 200 ("AI") structure according to the present invention. Generally, an Access Interval is the basic communication unit, a fixed block of time, that allocates bandwidth to synchronization, media access, polled

- 29 -

communications, contention based communications, and scheduled services. The Access Interval in Fig. 2 includes a SYNC header 201 generated by a Control Point ("CP") device of a NET. The term NET describes a group of users of a given hopping sequence or a hopping sequence itself. The Control Point device is generally the base station 15 referenced above with regard to Fig. 1. The SYNC header 201 is used by constituents of the NET to attain and maintain hopping synchronization. A reservation phase 203 follows permitting a reservation poll, which provides the NET constituents an opportunity to gain access to media. A sessions frame 205 is next allocated for communication protocol. A frame 207 follows for optional time division multiple access ("TDMA") slots in order to accommodate scheduled services. Scheduled services, for example, real time voice or slow scan video, are such that they require a dedicated time slot to provide acceptable quality of service. The function of frames 201, 203, 205 and 207 will be discussed in greater detail below.

As was shown in Fig. 2, Fig. 21 illustrates a sequence in an access interval 2100 with the Time Division Multiple Access slots 2113 positioned at the end of the access interval 2100. In present example, if this were also a HELLO interval, the HELLO would immediately follow the SYNC 1201. Location of the Time Division Multiple Access slots at such a position provides certain advantages including, for example, 1) the SYNC 2101, HELLO (not shown), Reservation Poll 2103, may all be combined into a single transmission (concatenated frames); 2) hopping information may be moved to or included in the Reservation Poll 2103 allowing for a shorter preamble in the SYNC 2101; and 3) the HELLO messages will occur early in the Access Interval 2100

- 30 -

providing for shorter receiver on times for sleeping terminals.

The Time Division Multiple Access slots may also be located at different points within the access interval. Positioning the Time Division Multiple Access slots allow for various systemic advantages. Referring now to Fig. 22, an access interval 2200 is illustrated showing the Time Division Multiple Access slots 2203 immediately following the SYNC 2201. Location of the Time Division Multiple Access slots 2203 at this position provides certain advantages including, for example, 1) better timing accuracy is achieved when the Time Division Multiple Access slots 2203 immediately follow the SYNC 2201; 2) Session Overruns do not interfere with the Time Division Multiple Access slots 2203; 3) devices which do not use the Time Division Multiple Access slots 2203 do not necessarily need to be informed of the Time Division Multiple Access slot allocation; and 4) HELLO message may follow Time Division Multiple Access slots 2203, Reservation Slots 2207 or Reservation Resolution Poll 2209.

Referring now to Fig. 23, an access interval 2300 is illustrated showing the Time Division Multiple Access slots 2305 immediately following the SYNC 2301 and the Reservation Poll 2303. In the present example, if this were a HELLO interval, a HELLO message would immediately follow the Reservation Resolution Poll 2309.

Location of the Time Division Multiple Access slots 2305 at the position shown in Fig. 23 provides certain advantages including, for example, 1) the Time Division Multiple Access slot timing is keyed to SYNC 2301 for better accuracy; 2) the number of Time Division Multiple Access slots 2305 may be indicated in SYNC 2301 or the Reservation Poll 2303, providing greater flexibility; 3) Session frame overruns do not interfere with Time

- 31 -

Division Multiple Access slots 2305; 4) only one maintenance transmission is required per Access Interval 2300; and 5) hopping information may be moved to or included in the Reservation Poll 2303, permitting a shorter preamble in SYNC 2301.

In the access interval 2300 configuration shown in Fig. 23, it is possible that the Time Division Multiple Access slots 2305 and the response slots 2307 could be the same. The Reservation Poll 2303 would allocate the correct number of slots and indicate which are reserved for Time Division Multiple Access. For example, to use Idle Sense Multiple Access 1 slot) with 1 inbound and 1 outbound Time Division Multiple Access slots, three slots would be allocated with the first two slots reserved. The appropriate Time Division Multiple Access slot duration is 80 bits at a hop rate of 200 hops per second which is just about the expected duration of a Request for Poll. At slower hop rates, multiple slots could be allocated to Time Division Multiple Access allowing the Time Division Multiple Access slot duration to be constant regardless of hop rate.

Referring now to Fig. 24, another access interval 2400 is illustrated showing the Time Division Multiple Access slots 2403 immediately following the SYNC 2401. In this example the Poll Message Queue 2405 immediately follows the Time Division Multiple Access slots 2403. The configuration shown in Fig. 24 provides for certain advantages including, for example, 1) the Time Division Multiple Access slot timing is keyed to SYNC 2401 for better accuracy; and 2) Session frame overruns do not interfere with Time Division Multiple Access slots 2403.

The configurations shown in Fig. 21 and in Fig. 23 are preferred because they allow the Reservation Poll messages to be transmitted immediately following the SYNC

- 32 -

and because of the power management and interference reduction advantages.

5 In one embodiment of the Access Interval structure, all message transmissions use standard high-level data link control ("HDLC") data framing. Each message is delimited by High-Level Data Link Control Flags, consisting of the binary string 01111110, at the beginning of the message. A preamble, consisting of a known data pattern, precedes the initial FLAG. This preamble is used to attain clock and bit synchronization prior to start of data. Receiver antenna selection is also made during the preamble for antenna diversity. A CRC for error detection immediately precedes the ending FLAG. Data is NRZ-I (differentially) encoded to improve data clock recovery. High-Level Data Link Control NRZ-I data is run-length-limited to six consecutive bits of the same state. Alternatively, a shift register scrambler could be applied instead of differential encoding to obtain sufficient transitions for clock recovery. Data frames may be concatenated, with two or more frames sent during the same transmission, with a single FLAG separating them. An example of this is SYNC, followed by a HELLO or Reservation Poll (SYNC, HELLO and Reservation Poll are discussed more fully below).

25 While much of the following discussion centers on the use of frequency hopping in the premises LAN, the Access Interval structure of the present invention is also suitable for single channel and direct sequence spread spectrum systems. The consistent timing of channel access, and the relative freedom from collisions due to channel contention, provide desirable benefits in systems that support portable, battery powered devices regardless of modulation type or channelization. Functions that are unique to frequency hopping may be omitted if other channelization approaches are used.

- 33 -

Figs. 3a and 3b illustrate the frequency of operation periodically changing corresponding to Access Interval boundaries in a frequency hopping system. Frequency hopping systems use a hopping sequence, which is a repeating list of frequencies of length (n) selected in a pseudo random order and is known to all devices within a coverage area. Fig. 3a illustrates a frequency hopping system having one Access Interval 301 per frequency hop (the hop occurring every 10 milliseconds) and a length of 79. Fig. 3b illustrates a frequency hopping system having one Access Interval 303 per frequency hop (the hop occurring every 20 milliseconds) and a length of 79. The 20ms time frame is preferred for a protocol stack that uses a maximum network layer frame of up to 1536 bytes payload while maintaining two real time voice communications channels. Access interval duration may be optimized for other conditions. Access Interval length is communicated to the NET during the SYNC portion of the Access Interval. This allows Access Interval duration, and other NET parameters to be adjusted without reprogramming every device within the NET.

The Access Interval is a building block. The length of the Access Interval can be optimized based on network layer packet size, expected mix of Bandwidth on Demand ("BWOD") and Scheduled Access traffic, expected velocities of devices within the NET, acceptable duration of channel outages, latency or delay for scheduled services, etc. The preferred Access Interval duration of 20ms (and maximum packet length of 256 Bytes at 1 MBIT/sec) represents a value chosen for systems with device velocities up to 15 MPH, and a mix between Bandwidth On Demand and scheduled service traffic.

Within a frequency hopping network, one or more Access Intervals may be used during each dwell in a

- 34 -

frequency hopping system. A dwell is the length of time (d) each frequency in the hopping sequence is occupied by the system. For example, Figs. 4a and 4b show illustrations of cases where more than one 20 ms Access Interval 401 is used per hop. This may be appropriate for some instances where it is undesirable to hop at higher rates because of relatively long frequency switching times of the radio hardware, where import, export, or regulatory restrictions disallow hopping at a faster rate, or in some applications where it is desirable to maintain operation on each channel for a longer period. An example of the latter is the case where larger files or data records are transferred routinely.

In a frequency hopping operation, the Access Interval 200 of Fig. 2 begins with a SYNC header 201. As mentioned above, the SYNC is generated by the Control Point (CP) device of the NET. The SYNC is used by constituents of the NET to attain and maintain hopping synchronization. Included in the SYNC are:

1. Address of the Control Point device.
2. Identification of the Hopping Sequence, and index of the current frequency within the hop table.
3. Identification of the hop rate, number of Access Intervals per hop, and Access Intervals before next hop.
4. A timing character for synchronization of device local clocks to the NET clock contained within the Control Point device.
5. Status field indicating reduced SYNC transmissions due to low NET activity (Priority SYNC Indicator).
6. Status field indicating if the Access Interval will contain a broadcast message to all devices within the NET.

- 35 -

7. Status field indicating premises or spontaneous LAN operation.
8. The SYNC field information is optionally encrypted using a block encryption algorithm, with a key provided by the network user. A random character is added to each SYNC message to provide scrambling.

However, there are two circumstances during which a SYNC message is not transmitted: 1) co-channel interference; and 2) low NET utilization. With regard to co-channel interference, before issuing a SYNC message, the Control Point device performs channel monitoring for a brief interval. If the Received Signal Strength Indicator (RSSI) level indicates an ON channel signal greater than the system defer threshold, then the Access Interval is skipped. Alternatively, a strong ON channel signal may dictate a reduction in Control Point device power to limit the interference distance of the net for the duration of the Access Interval. A system defer threshold 30 dB above the receiver sensitivity is a preferred choice. Communication within the NET is deferred for the duration of the Access Interval if SYNC is not transmitted due to co-channel interference.

In times of low system utilization, SYNC and Reservation Poll messages are reduced to every third Access Interval. The SYNC message includes a status field indicating this mode of operation. This allows devices to access the NET, even during Access Intervals where SYNC is skipped, by using an Implicit Idle Sense

- 36 -

algorithm. If the hopping sequence is 79 frequencies in length as shown in Figs. 3a and 3b, use of every third Access Interval guarantees that a SYNC message will be transmitted on each frequency within the hopping sequence once each three cycles of the sequence, regardless of whether 1, 2 or 4 Access Intervals occur each hop dwell. This addresses US and European regulatory requirements for uniform channel occupancy, and improves the prospects for synchronization of new units coming into the NET during periods when the NET is otherwise inactive. SYNC messages that are on multiples of 3 Access intervals are labeled as priority SYNC messages. "Sleeping" terminals use priority SYNCs to manage their internal sleep algorithms. Sleeping terminals and Implicit Idle Sense are discussed in more detail below.

It should be noted that SYNC messages are preceded by dead time, which must be allocated to account for timing uncertainty between NET clocks and local clocks within NET constituents. In frequency hopping systems, the dead time must also include frequency switching time for the RF modem.

The Reservation Poll frame 203 immediately follows the SYNC header 201. The two messages are concatenated High-Level Data Link Control frames separated by one or more Flags. The reservation poll provides NET constituents an opportunity to gain access to the media. It includes:

- 37 -

1. A field specifying one or more access slots.
2. A field specifying a probability factor between 0 and 1.
3. A list of addresses for which the base stations has pending messages in queue.
4. Allocation of Time Division Multiple Access slots for scheduled services by address.
5. Control Point device Transmitted Power level for SYNC and Reservation Polls.

The number of access slots, n , and the access probability factor, p , are used by the Control Point device to manage contention on the channel. They may each be increased or decreased from Access Interval to Access Interval to optimize access opportunity versus overhead.

If the NET is lightly loaded, the pending message list is short, and the NET is not subject to significant interference from other nearby NETs, the control point device will generally specify a single slot 501 as shown in Fig. 5a, with a p factor < 1 . In this case, the reservation phase is Idle Sense Multiple Access ("ISMA"). Devices with transmission requirements that successfully detect the Reservation Poll will transmit a Request for Poll ("RFP") with probability p and defer transmission with probability $1-p$. Fig. 5b shows a device response (address 65 503 following the reservation poll).

In cases when the transmission density is higher, n multiple reservation slots will be specified, generally with a probability factor p of 1. In this case a device

- 38 -

will randomly choose one of n slots for transmission of their Request for Poll. The slotted reservation approach is particularly appropriate in instances where many NETs are operating in near proximity, since it diminishes
5 reliance on listen before talk ("LBT") (explained more fully below). The number of slots n is determined by a slot allocation algorithm that allocates additional slots as system loading increases. Fig. 6a shows multiple slots 601.

10 In cases where NET loading is extreme, the Control Point may indicate a number of slots, e.g., not more than 6, and a probability less than 1. This will cause some number of devices to defer responding with a Request for Poll in any of the slots. This prevents the control
15 point device from introducing the overhead of a large number of slots in response to heavy demand for communications, by dictating that some units back off until demand diminishes.

A pending message list is included in the
20 Reservation Poll. The pending message list includes the addresses of devices for which the Control Point device has messages in queue. Devices receiving their address may contend for the channel by responding with a Request For Poll (RFP) in the slot response phase. Fig. 6b shows
25 several devices 603, 605 and 607 contending for channel access. Messages that the Control Point device receives through the wired infrastructure that are destined for

- 39 -

Type 1 devices, and inactive Type 3 devices whose awake window has expired, are immediately buffered, and the device addresses are added to the pending message list. When a message is received through the infrastructure for a Type 2 device, or an active Type 3 device, their address is prioritized at the top of the polling queue. (Device Types and polling queue are described below.) The pending message list is aged over a period of several seconds. If pending messages are not accessed within this period, they are dropped.

Devices with transmission requirements respond in slots with a Request for Poll. This message type includes the addresses of the Control Point device and requesting device, the type and length of the message it has to transmit, and a field that identifies the type of device. Devices that detect their address in the pending message list also contend for access in this manner.

As mentioned above, devices may be Type 1, Type 2, or Type 3. Type 1 devices are those which require critical battery management. These may be in a power saving, non-operational mode much of the time, only occasionally "waking" to receive sufficient numbers of SYNC and Reservation Poll messages to maintain connectivity to the NET. Type 2 devices are those that are typically powered up and monitoring the NET at all times. Type 3 units are devices that will remain awake for a window period following their last transmission in

- 40 -

anticipation of a response. Other device types employing different power management schemes may be added.

Slot responses are subject to collision in both the single and multiple slot cases. Collisions may occur when two or more devices attempt to send Request for Polls in the same slot. However, if the signal strength of one device is significantly stronger than the others, it is likely to capture the slot, and be serviced as if it were the only responding unit. Fig. 6b shows two devices 605, address 111, and 607, address 02, that may be subject to collision or capture.

The Control Point device may or may not be able to detect collisions by detecting evidence of recovered clock or data in a slot, or by detecting an increase in RF energy in the receiver (using the Received Signal Strength Indicator, ("RSSI")) corresponding to the slot interval. Collision detection is used in the slot allocation algorithm for determining addition or deletion of slots in upcoming Reservation Polls.

As an optional feature to improve collision detection in the multiple slot case, devices that respond in later slots may transmit the addresses of devices they detect in earlier slots as part of their Request for Poll. Request for Polls which result in collisions at the Control Point device often are captured at other remote devices, since the spatial relationship between devices that created the collision at the base does not

- 41 -

exist for other device locations within the NET. The duration of the response slots must be increased slightly to provide this capability.

5 If the Control Point device receives one or more valid Request for Polls following a Reservation Poll, it issues a Reservation Resolution ("RR") Poll and places the addresses of the identified devices in a polling queue. The Reservation Resolution message also serves as a poll of the first unit in the queue. Addresses from
10 previous Access Intervals and addresses of intended recipients of outbound messages are also in the queue.

If the Polling Queue is empty, then no valid Request for Polls were received or collision detected and no Reservation Resolution poll is issued. If within this
15 scenario a collision is detected, a CLEAR message indicating an Explicit Idle Sense (explained more fully below) is transmitted containing a reduced probability factor to allow colliding units to immediately reattempt NET access.

20 Outbound messages obtained through the network infrastructure may result in recipient addresses being prioritized in the queue, that is, if the recipients are active devices -- Type 2 devices or Type 3 devices whose awake window has not expired. This eliminates the need
25 for channel contention for many outbound messages, improving efficiency. Messages for Type 1 devices are

- 42 -

buffered, and the recipient address is placed in the pending message list for the next Access Interval.

Generally the queue is polled on a first in first out (FIFO) basis. The polling order is:

- 5 a. Addresses of active units with outbound messages.
- b. Addresses from previous Access Intervals
- c. Addresses from the current Access Interval

10 Since propagation characteristics vary with time and operating frequency, it is counterproductive to attempt retries if Poll responses are not received. If a response to a Poll is not received, the next address in the queue is polled after a short response time-out period. Addresses of unsuccessful Polls remain in the

15 queue for Polling during the next Access Interval. Addresses are aged, so that after several unsuccessful Polls they are dropped from the queue. Addresses linked to outbound messages are added to the pending message list. Devices with inbound requirements must re-enter

20 the queue through the next reservation phase.

 Data is transferred in fragments. A maximum fragment payload of 256 bytes is used in the preferred implementation. If transfer of network packets larger than of 256 bytes is required, two or more fragments are

25 transferred. Fragments may be any length up to the maximum, eliminating the inefficiency that results when messages that are not integer multiples of the fragment

- 43 -

length are transmitted in systems that employ fixed sizes.

The sequence for transferring data from a remote device to the control point device is illustrated in Fig. 7a. It is assumed that address 65 is the first address in the polling queue. The Reservation Resolution poll 701 from the control point device includes the device address and the message length that device 65 provided in its initial Request for Poll. A first fragment 703 transmitted back from device 65 is a full length fragment. Its header includes a fragment identifier and a field providing indication of the total length of the message. Length information is included in most message types during the sessions period to provide reservation information to devices that may wish to attempt to access the NET following an Explicit Idle Sense (explained more fully below).

Following successful receipt of the first fragment, the Control Point device sends a second poll 705, which both acknowledges the first fragment, and initiates transmission of the second. The length parameter is decremented to reflect that the time required for completion of the message transfer is reduced. A second fragment 707 is transmitted in response, and also contains a decremented length field. Following receipt of the second fragment 707, the Control Point device sends a third poll 709. This pattern is continued until

- 44 -

a final fragment 711 containing an End of Data (EOD) indication is received. In Fig. 7, the final fragment is shorter than a maximum length fragment. The Control Point device sends a final Acknowledge (ACK), and the device sends a final CLEAR 713 to indicate conclusion of the transmission. The CLEAR message contains a probability factor p for Explicit Idle Sense (explained more fully below). The value of p is determined by the Control Point device in the ACK and echoed by the device termination communication. A p of zero indicates that the control point device will be initiating other communications immediately following receipt of the CLEAR message. A probability other than 0 indicates an Explicit Idle Sense.

If for some reason a fragment is not successfully received, the next poll from the Control Point device would indicate a REJECT, and request re-transmission of the same fragment. The length field would remain fixed at the previous value, prolonging reservation of the channel for the duration of the message. After a fragment is transmitted more than once without successful reception, the Control Point device may suspend attempts to communicate with the device based upon a retry limit, and begin polling of the next address in the queue.

A flow chart depicting how inbound messages are received during an access interval is shown in Figs. 19a, 19b-1 and 19b-2. A flow chart depicting how outbound

- 45 -

messages are transmitted during an access interval is shown in Figs. 20a, 20b-1 and 20b-2.

Outbound messages are transmitted in a similar fashion as inbound messages, with the Control Point and device roles largely reversed as illustrated in Fig. 7b. When the Control Point reaches an address in the queue for which it has an outbound message, the Control Point transmits a Request for Poll 721 identifying the address of the device and the length of the message. The response back from the device would be a poll with an embedded length field. The same POLL/FRAGMENT/ACK/CLEAR structure and retry mechanisms as described above with regard to inbound messages in reference to Fig. 7a are maintained. The CLEAR from the device indicates a probability p of zero. If the polling queue is empty, the Control Point may send a final or terminating CLEAR 723 containing a probability for Explicit Idle Sense.

All terminating ACK or CLEAR messages contain fields to aid in synchronization of new units to the NET. The content of these fields is identical to that in the SYNC message, except that the timing character is deleted. Synchronization is discussed more fully below.

Broadcast Messages intended for groups of addresses, or all addresses within a NET may be transmitted during the sessions period. Broadcast messages are not individually acknowledged. These messages may be communicated at intervals over the course of several

- 46 -

Access Intervals to provide reliable communication. Messages such as SYNC and Reservation Polls are specialized broadcast messages, with dedicated bandwidth in the Access Interval structure.

5 Security of payload data is left to the higher protocol layers. Application programs resident in portable/mobile devices may employ encryption or other means of providing protection against undesired use of transmitted data.

10 Portable/mobile devices may employ transmitter power control during the sessions period to reduce potential interference with other NETs that may occasionally be on the same or adjacent channels. These devices will use Received Signal Strength Indicator readings from outbound
15 messages to determine if transmitter power may be reduced for their inbound transmission. Because of the need to maintain channel reservations and Listen Before Talk capabilities, the Control Point device does not use transmitter power control. Since Control Point devices
20 are generally part of an installed system infrastructure, they are likely to be physically separated from devices operating in other NETs. They are therefore less likely to cause interference to devices in other NETs than portable devices, which may operate in proximity to
25 devices in other NETs.

 Often, control point devices will empty the polling queue before the conclusion of the access interval. Two

- 47 -

mechanisms within the Access Control Protocol, Explicit and Implicit Idle Sense, are provided to improve bandwidth utilization. These supplemental access mechanisms often provide means for devices that failed to gain reservations during the reservation phase to gain access to the NET within the Access Interval. To assume an Explicit or Implicit Idle Sense, a device must have detected a valid SYNC and Reservation Poll in the current Access Interval.

The incorporation of a probability factor $p \neq 0$ in the final (terminating) ACK or CLEAR from the control point device provides the function of an Explicit Idle Sense (mentioned above). Devices with transmission requirements solicit Request for Polls using the same rules normally used for a single slot Reservation Poll. Successfully identified addresses are placed in the polling queue, and are polled immediately or in the next Access Interval depending on the time remaining in the current Access Interval. The p factor for Explicit Idle Sense is subject to the same optimization algorithm as the Reservation Poll probability.

Communication of channel reservations, in the form of the length fields in Polls and Message Fragments is useful to units seeking to access the NET through Explicit Idle Sense. Reservations allow devices to predictably power down during the period that another

- 48 -

device has reserved the NET to conserve battery power, without losing the ability to gain access to the NET.

Implicit Idle Sense provides an additional means of channel access. An Implicit Idle Sense is assumed
5 whenever a device detects a quiet interval period greater than or equal to the duration of a Poll plus the maximum fragment length after a channel reservation has expired. Detection based upon simple physical metrics, such as a change in Received Signal Strength Indicator or lack of
10 receiver clock recovery during the quiet interval, are preferred methods of ascertaining channel activity. Algorithms based upon these types of indicators are generally less likely to provide a false indication of an inactive channel than those that require successful
15 decoding of transmissions to determine channel activity. False invocation of an Implicit Idle Sense is the only mechanism by which data transmissions are subject to collision within the NET. Thus, the Implicit Algorithm must be conservative.

20 Quiet interval sensing may begin at the following times within the Access Interval:

- a. Any time after the last reservation slot following a Reservation Poll;
- 25 b. Any time after a terminating ACK or CLEAR indicating an Explicit Idle Sense;
- c. Following an unsuccessful response to a single Slot Reservation Poll; or

- 49 -

- d. Any time prior to reserved Time Division Multiple Access time slots at the end of the Access Interval.

5 It is preferable that devices detecting a quiet interval use a *p* persistent algorithm for channel access to avoid collisions. The probability factor for Implicit Idle Sense Access will generally be less than or equal to the factor in Explicit Idle Sense.

10 A device must receive the SYNC and Reservation Polls at the beginning of an Access Interval to use Implicit Idle Sense. The Reservation Poll provides indication of guaranteed bandwidth allocation to scheduled services at the end of the Access Interval, which may shorten the period available for Bandwidth On Demand communications.

15 Devices requiring scheduled services must contend for the channel in the same fashion as those requiring Bandwidth On Demand access. When polled, these initiating devices will initiate a connection request that indicates the number of inbound and outbound Time
20 Division Multiple Access slots required for communication, and the address of the target device with which communication is desired. The network infrastructure will then attempt to establish the connection to the target device. Once the connection is
25 established, the Control Point device will signal the allocation of slots to the initiating device. Time Division Multiple Access slots are relinquished by transmitting a disconnect message to the control point

- 50 -

device in the Time Division Multiple Access slot until the disconnect is confirmed in the next Reservation Poll.

The transmission requirements of speech and slow scan video (scheduled services) are similar. In one
5 embodiment, Time Division Multiple Access slots are allocated as multiples of 160 bits payload at 1 MBIT/sec, plus overhead for a total of 300 μ s. For 10 ms access intervals, acceptable voice communication can be obtained by allocating 1 Time Division Multiple Access slot each
10 for inbound and outbound communication per access interval. For 20 ms access intervals, two slots each way are required. A system employing 10 ms access intervals at 100 hops per second may improve transmission quality by using two or three slots each Access Interval and
15 sending information redundantly over two or three access intervals using interleaved block codes. Scheduled transmissions are generally not subject to processing or validation by the control point device, and are passed through from source to destination. Use of interleaved
20 error correction coding or other measures to improve reliability are transparent to the NET.

The selection of certain system parameters are important when considering scheduled services. As an
example, since speech is quantized over the duration of
25 the access interval and transmitted as a burst, the length of the access interval translates directly into a transport delay perceptible to the recipient of that

- 51 -

speech. In real time voice communications, delays longer than 20 ms are perceptible, and delays longer than 30 ms may be unacceptable. This is particularly the case where the premises LAN is interconnected with the public switched telephone network ("PSTN"), which introduces its own delays. Two way services such as voice communications are the most sensitive to transport delay because delay impacts the interaction of the communicating parties. One way services are less sensitive to transport delay. One way services are good candidates for interleaving or other forms of redundant transmission.

Similarly, the selection of hop rate is important, as hop rate determines the duration of outages that may occur. If one or more frequencies in the hop sequence are subject to interference, for instance, scheduled transmissions during those hops will be disrupted. In a system that hops slowly, detrimental outages of hundreds of milliseconds will occur resulting in poor transmission quality. Occasional losses of smaller durations, e.g., 10 ms or 20 ms, are generally less perceptible, indicating that faster hop rates are desirable if the NET is to offer real time voice transport.

Scheduled service intervals may also be used for data transport on a scheduled or priority basis. Telemetry, data logging, print spooling, modem replacement, or other functions are possible. For these

- 52 -

activities, a few Time Division Multiple Access slots scheduled for example every fourth, eighth, or sixteenth Al are necessary.

5 Because of multipath and dispersion issues with 2.4 GHz transmission at relatively high data rates, the ability of the NET to adaptively switch between two or more data rates is desirable.

10 In one embodiment, implementation of data rate switching may be accomplished by selecting a standard rate of communications, e.g., 250 KBPS and high rate of communications of 1 Mbit/sec. Messages that contain system status information, including SYNC, Reservation Polls, Reservation Resolution Polls (Request for Polls), Polls, ACKs and CLEARS are transmitted at the standard
15 rate. These messages are generally short, and the time required for transmission is largely determined by hardware overhead, e.g., transmitter receiver switching time. The incremental overhead introduced by transmitting these messages at the lower rate is
20 therefore small in comparison to the total length of an access interval. The reliability of reception of these messages will increase, which will eliminate unnecessary retries in some instances where fragments are received successfully, but acknowledgements or polls are missed.

25 A test pattern at the higher data rate is inserted in each Poll (not in Reservation Polls, however). The Poll recipient evaluates signal quality based on the high

- 53 -

data rate test pattern, Received Signal Strength Indicator, and other parameters to determine whether to transmit a fragment at the high rate or the low rate. Fragment lengths are selected such that high and low rate maximum fragment lengths are the same duration. In other words, a fragment at the low rate conveys approximately 1/4 the payload of a fragment for the case where the data rate is four time greater. This method is generally suitable for transaction oriented communications, which frequently require short message transmissions. Alternatively, the length field in Polls and messages can be used to allow different fragment lengths for the two data rates while still providing channel reservation information to other devices in the NET. This method also provides for forward migration. As modulation and demodulation methods improve, newer products can be added to old networks by upgrading Control Points devices. Both new and old devices share the ability to communicate at a common low data rate.

An alternate embodiment uses signaling messages such as SYNC, Reservation Polls, Request for Polls, etc., at the higher rate with fallback operation to the standard rate for the communications sessions only. SYNC and Reservation Polls at the high rate constitute a high data rate test message. The Request for Poll response to the Reservation Poll at the high rate may include a field indicating that sessions communications should take place

- 54 -

at the fallback, standard rate. Signal quality measures such as signal strength and clock jitter are appropriate. Data rate selection information is included with the device address in the polling queue. When the device is
5 polled, it will be polled at the rate indicated in the Request for Poll. Channel reservation information in the Reservation Resolution Poll will indicate the reservation duration based upon the data rate indicated.

In this alternate embodiment, the fact that SYNC and
10 Reservation Polls must be detectable at the high data rate prioritizes access to the NET for those devices that have acceptable connectivity during the current access interval. This general approach has desirable characteristics in a frequency hopping system, as the
15 propagation characteristics between devices may change significantly as the NET changes from frequency to frequency within the hopping sequence, or over several Access Intervals during the dwell time on a single frequency. Reduction in data rate in this system is
20 primarily intended to remedy the data smearing (inter-symbol interference) effects of dispersion due to excess delay, rather than temporary poor signal to noise ratio due to frequency selective fading. Devices that receive high data rate transmissions with acceptable signal
25 strength but high jitter are likely to be experiencing the effect of dispersion.

- 55 -

The concept of allowing Polls and message fragments to occur at either a high or low data rate could create difficulties for other NET constituents that need to be able to monitor the channel for reservation information.

5 Two embodiments for solving this problem are the use of auto-discriminating receivers or the use of fixed data rate headers for system communications.

Auto discrimination requires the receiver to process messages sent at either data rate, without necessarily
10 having prior knowledge of the rate.

Given a high rate of 1 MBIT/SEC, and a low Rate of 250 Kbps, i.e., one being a binary multiple of the other, it is possible to devise preambles that can be received at either rate. Consider that 01 and 110 sent at the low
15 rate correspond to 00001111 and 111111110000 at the high rate. These preambles are transmitted continuously before the transmission of the High-Level Data Link Control FLAG character at the correct data rate indicating the start of a message. In this example, a
20 preamble of 20 bits of 01 at the low rate indicates operation at the high rate. A preamble of 30 bits of 110 indicates operation at the low rate. A receiver tuned to either rate is capable of receiving both types of preambles and initiating the proper decoding mechanisms
25 for the intended rate of transmission.

This general technique, with appropriate selection of preamble content, is applicable to binary modulation

- 56 -

schemes, for example, a frequency modulated system where a common frequency deviation value is used for both data rates. It is also applicable to systems where switching occurs between binary and multilevel modulation, such as disclosed in pending U.S. Application Ser. No. 07/910,865, filed July 6, 1992.

Referring now to Fig. 25, a preamble 2501, a SYNC 2503 and a Reservation Poll 2505 is illustrated. The preamble 2501 starts at the beginning of the Access Interval 2500 and is applied to an RF modem while it is switching frequencies. Since the switching time is a worst case, this causes the preamble 2501 to be present and detectable prior to the allocated 150 μ sec period in some instances. It would be equally appropriate to begin preamble transmission 50 or 100 μ sec into the switching period if that would be more convenient. The timing has been selected to allow 100 μ sec.

Referring to Fig. 26, a sample SYNC message 2600 is shown. Referring to Fig. 27, a sample Reservation Poll 2700 is shown. In these examples, the hopping synchronization information has been positioned in the Reservation Poll 2700.

With auto-discrimination, it is possible to change data rates on a per-poll basis, thereby adjusting for channel temporal dynamics. Since all devices in the NET have auto discrimination capabilities, and channel reservation information is included in message headers as

- 57 -

a length field, the bandwidth reservation features of the NET are preserved. The maximum fragment duration may be maintained at a fixed value, meaning that low data rate fragments convey less data than their high rate counterparts, or may be scaled in the ratio of the data rates to allow consistent fragment data payloads.

An alternative to auto-discrimination is the use of headers to communicate system information. This embodiment is less preferred, but may be appropriate if economics, size, or power constraints dictate a simpler design than that required for auto-discrimination. In this embodiment, any transmission at the lower data rate is preceded by a header at the high data rate that conveys NET management information, i.e., channel reservation status. Devices other than those directly involved in polling or fragment transmission need only monitor at the high rate for channel reservation information. The header at the high rate and the following transmission at the low rate are concatenated High-Level Data Link Control frames, with an appropriate preamble for low rate clock recovery synchronization in-between.

For the communicating devices, the header can serve the additional purpose of acting as a test pattern at the high rate. For example, if a device is polled at the low rate, but successfully decodes the high rate header with

- 58 -

adequate signal quality, it may indicate back to the polling unit to poll again at the high rate.

In a premises LAN as discussed in reference to Fig. 1, many NETs may be distributed geographically to provide enhanced coverage or additional system capacity. The
5 wired portion of the network infrastructure, such as Ethernet or Token Ring, provides a means for coordination of NETs to achieve optimum system performance. An equally important role of the wired infrastructure is to
10 allow resource sharing. Portable devices with limited memory capacities, processing power, and relatively small batteries may access large data bases on, or remotely initiate processing capabilities of, larger AC powered computer systems. Portable/mobile devices may also share
15 communication with other like devices which are serviced by other NETs well beyond the radio coverage range of their own NET.

The basic method for communication of status information regarding the premises LAN is the HELLO
20 message. HELLO messages are sent routinely, but relatively infrequently, for example, every 90 Access Intervals. The HELLO transmission interval is tied to the Priority SYNC interval, so that the HELLO interval corresponds to Access Intervals where SYNC is transmitted
25 if the network is lightly utilized.

In an alternate embodiment, HELLOs could be inserted as a broadcast message at the beginning of the Sessions

- 59 -

period. Fig. 8 illustrates a preferred Access Interval embodiment where a HELLO message 801 is inserted between a SYNC 803 and a Reservation Poll 805. The SYNC frame at the beginning of the Access Interval indicates that the
5 Access Interval will contain a HELLO, allowing power managed devices to remain awake to receive the HELLO.

HELLO messages may also contain information regarding pending changes in the local NET. If the local NET is changing Access Interval durations or hop
10 sequences, for instance, changes may be communicated in several consecutive HELLOs so that the information is reliably communicated to all NET constituents, permitting all devices to make the change in coordinated fashion. Further discussion of HELLO message content is provided
15 below.

For purposes of channel management in the Access Interval structure, the maximum transmission duration by a device should be limited to the time that the device moving at a maximum expected velocity can traverse 1/4
20 wavelength of the maximum carrier frequency. The duration may be further reduced to compensate for link bit error rate characteristics or expected duration or frequency of interference bursts. A maximum transmission duration of 2.5 ms is suitable for 1 MBIT/SEC
25 transmission, with a device velocity of 15 mph, in a multiple NET environment.

- 60 -

Use of spatial or polarization antenna selection diversity is also desirable in indoor propagation environments. First, the receiving unit makes an antenna diversity decision during the preamble portion of each transmission. The antenna used for reception for each device address is then recorded in memory so that the correct antenna will be used for response messages to each address. While diversity selection is only valid for a short time, it is not necessary to age this information, because antenna selection is equi-probable even after diversity information is no longer valid.

The Access Interval structure of the present invention also inherently provides routine channel sounding for each hop. This is important in a frequency hopping system, as channel conditions will vary considerably from frequency to frequency within the hopping sequence. NET constituents must, in most cases, be able to receive SYNC and Reservation Poll transmissions from the Control Point device to attempt inbound access in an Access Interval. This provides a positive indication that the device is not experiencing a channel outage, allowing power saving and eliminating possible channel contention. Channel sounding does not need to be employed during periods where the NET is not busy since contention is unlikely in this situation.

Channel sounding for Outbound messages is accomplished through a Request for Poll/Poll cycle where

- 61 -

handshaking messages with short time out periods must be successfully communicated before longer message transmissions may be attempted.

As discussed above with regard to Fig. 1, a premises
5 LAN consists of several base stations 15 located
throughout an environment requiring wireless
communications, e.g., a building or other facility, or a
campus comprised of several buildings. The base stations
15 are placed to provide coverage of intended usage areas
10 for the roaming portable or mobile computing devices 20.
Coverage areas must overlap to eliminate dead spots
between coverage areas.

The base stations 15 may be interconnected via
industry standard wired LANs, such as IEEE 802.3
15 Ethernet, or IEEE 802.5 Token Ring. Base stations may be
added to an existing LAN without the need to install
additional LAN cable. Alternatively, it may be desirable
to install base stations on dedicated LAN segments to
maximize performance of both the radio network and other
20 collocated computer devices.

Base stations within the premises LAN provide
Control Point functions for individual NETs. NETs employ
different hopping sequences to minimize potential
interference between NETs. Regulatory restrictions
25 generally preclude synchronization of multiple NETs to a
single master clock, requiring that individual NETs
operate independently from one another. The lack of the

- 62 -

ability to coordinate timing or frequency usage between NETs introduces the potential for collisions between independent NETs with overlapping coverage areas.

5 Figs. 9a and 9b illustrate conceptually how multiple NETs may be employed in an idealized "cellular" type installation. Each hexagon 901 and 903 in Fig. 9a represents the primary coverage area of a given NET. Coverage areas are modeled as circles 905 based upon some reliability criterion, for example a 5% mean fragment
10 retry rate (on average 95% of fragments are successfully communicated on the first attempt). Typical coverage areas are determined by physical attributes of the area in which the NET operates. As is illustrated in Fig. 9b for the hexagon (NET) 903 of Fig. 9a, an actual coverage
15 area 907 meeting the reliability criterion is likely to be irregular. This may require base stations to be offset significantly from the hexagonal grid.

Fig. 10 illustrates a coverage contour overlap for the multiple NETs in the premises LAN of Fig. 1. Darken
20 shaded areas 1001 indicate areas where base station coverage overlaps. Because the coverage distance of a radio system on an instantaneous basis greatly exceeds the coverage that can be provided on average to sustain a given quality of service, the overlap at any instant
25 may be significantly greater than the coverage contours indicate.

- 63 -

Fig. 11 illustrates hopping sequence reuse in a multiple NET configuration. Hopping sequence re-use may be necessary if there are physical constraints on the number of hopping sequences that can be supported. For example, devices may have limited memory available for hopping sequence storage. Use of a smaller set of sequences also simplifies the task of determining sets of sequences that have acceptable cross correlation properties. In Fig. 12, 7 hopping sequences 1 through 7 are used throughout the coverage area. Other NETs may reuse the same hopping sequence at some distance removed. While 7 NETs are illustrated, larger numbers, such as 9 or 15 may provide a better compromise between minimizing the number of hopping sequences used, and reuse distance between NETs using the same sequence. Reuse requires coordination of hopping sequence assignment -- either the system installer can coordinate the installation, or the system may include automated management features to assign hopping sequences to individual NETs.

Since NETs are not synchronized, different NETs that use the same hopping sequence are likely to interfere during periods where oscillator drift causes them to be temporarily synchronized. At other times, they may only interfere due to imperfect channelization. For example, for a worst case 100 ppm frequency error between two NETs using the same 79 frequency sequence at one Access Interval per hop and 50 hops per second, NETs will

- 64 -

partially or fully overlap for a duration of 10 minutes every 4.3 hours. Typically the frequency error will be 25% to 50% of the worst case, leading to longer overlap periods occurring less frequently.

5 NETs using the same hopping sequence must be physically isolated from one another to reduce interference to an acceptable level. Extensive hopping sequence reuse generally requires site engineering and optimization of base station placement. Using more
10 hopping sequences reduces the need for critical system engineering during installation. Fifteen hopping sequences is a preferred number for hopping sequence reuse, allowing simplified installation and minimal coordination.

15 NETs that use different hopping sequences will also temporarily synchronize in timing relationships that cause mutual co-channel interference on common channel frequencies. Since the number of channels that must be used in a sequence is a significant fraction of the total
20 number of channels available, all sequences will share some number of frequencies in common. When sequences are time aligned so that a common frequency is used simultaneously, interference can occur. Optimization of sets of sequences for low cross correlation is necessary
25 to prevent various time alignments of sequences from having more than one or two frequencies in common.

- 65 -

Optimization of hopping sequences for multiple NETs must also include analysis of imperfect channelization. The performance characteristics of the RF modems may not, for economic or power consumption reasons, provide sufficient transmitter spectral containment, receiver dynamic range, or receiver selectivity to guarantee that devices operating on different frequencies in proximity to one another will not interfere. In selecting hopping sequences for desirable cross correlation properties, adjacent and alternate adjacent channel interference must be considered. Protocol retry mechanisms for fragments lost to adjacent channel interference or limited dynamic range may be randomized to prevent continued disruption of communications in the affected NET.

Often in campus environments where systems must provide coverage in several buildings, the cost of wiring LAN cable between base stations is prohibitive. To establish connectivity between base stations in an premises LAN, it may be necessary to provide wireless links between groups of base stations connected to separate LAN segments. Fig. 12 illustrates a wireless link 1201 connecting groups of base stations 1203 and 1205. The base stations 1203 and 1205 are connected on separate LAN segments 1207 and 1209.

In one embodiment, the base stations 1203 and 1205 may be configured in a wireless point to point mode, wherein one base station serves as a control point device

- 66 -

while the others operate in a slaved mode dedicated to point to point data transfer. Slave base stations are configured to operate as portable/mobile devices, and forward communications to master bases by sending Request for Polls during reservation opportunities or Implicit Idle Sense periods. Because of the potential high traffic of point to point links, separate NETs may be allocated for this purpose, with a master communicating with one or more slave units. Master units may also communicate with other portable/mobile devices. The COST weighing (discussed below) in a slave's HELLO transmission is preferably set to a high value, to force portable/mobile devices which can connect to another NET to do so.

In another embodiment, it may also be desirable to support wireless base stations. Wireless base stations serve as control points, but are not connected to the infrastructure through a LAN cable. As is illustrated in Fig. 13, a wireless base station 1301 participates in the premises LAN through a wireless link 1303 to a base station 1305 that is connected to a LAN 1307.

Wireless base stations operate as slave devices to master base stations which are connected to the wired infrastructure. The wired and wireless base stations share the same hopping sequence, and are synchronized as a common NET. Because they are not connected to the Infrastructure, wireless base stations must be used as

- 67 -

store and forward devices. Each transmission to a wireless base must be retransmitted to the intended destination device, doubling the number of transmissions occurring in the NET. Wireless base stations are preferably used for supplementing coverage area of the premises LAN. For example, a wireless base station might provide spot coverage of isolated "dead spots" where data traffic is limited or where providing a wired LAN connection is difficult. Wireless base stations may also serve as emergency spares to provide coverage in the event of a failure of a primary base station. In this role, the wireless base station may be either permanently installed in selected locations, or stored in a maintenance area and quickly positioned and connected to AC or battery power to provide communications while repairs are made to the primary wired base station. Moreover, permanently installed wireless base stations might also be used for redundancy, i.e., to monitor an associated base station and to take over when a breakdown is detected.

The preferred wireless base station embodiment uses interleaved access intervals. The parent wired base station and secondary wireless base station coordinate Access Intervals, the wired base station deferring every third or sixth access interval to the wireless base. Since the wired base station transmits priority SYNC messages every third Access Interval, the wireless base

- 68 -

station may routinely be allocated one of the two
intervening Access Intervals for priority SYNC
communications with devices that are attached to it.
Communication between the wired and wireless base
stations may occur during Access Intervals initiated by
either base station. Wireless base stations may also
communicate with devices during an Access Interval using
Implicit or Explicit Idle Sense.

This embodiment provides predictable access for
devices attached to the wireless NET, and allows the same
power management algorithms to be used regardless of
whether the base station is wired or wireless. The
wireless base station may transmit its own priority SYNC
and HELLO messages. Also, devices seeking communications
with the wireless base station will automatically be
synchronized with the wired base as well, allowing
immediate improved access to the network if their
mobility has put them within range of the wired base.

Because of the constraint of sharing bandwidth with
a wired base station, connectivity of wireless base
stations is normally limited to one per wired base
station. However, in cases where system loading is
predictably and consistently light, multiple wireless
base stations could share a single wired base, e.g., each
transmitting in turn in the Access Intervals between the
Wired Base Priority SYNC Access Intervals.

- 69 -

Wireless base stations are capable of supporting scheduled traffic. However, since each transmission to a wireless base station must be forwarded, scheduled transmissions through wireless base stations use twice
5 the bandwidth as those through wired base stations. In other words, twice the number of Time Division Multiple Access slots must be allocated. To avoid introducing excessive delay, communications must be forwarded during the same Access Interval that they are received, or
10 shorter Access Intervals must be used. Scheduled traffic slot assignments must be common to all wireless bases operating within a single NET.

Wireless base stations require reliable communication with their wired counterparts. This
15 dictates smaller coverage contours for wireless base stations. If a wired base station provides 80,000 square feet of coverage area, a wireless base can be predicted to provide only an additional forty percent coverage improvement, due to overlap with the wired base station.
20 Frequently, base stations are mounted at ceiling level, providing a relatively clearer transmission path between base stations than exists between bases and portable/mobile devices located in more obstructed areas near the floor. With careful site engineering and
25 installation, a wireless base station can provide somewhat better than the forty percent predicted

- 70 -

improvement, but still less than the coverage of an additional wired base.

As discussed above, HELLO messages are used to communicate NET and premises LAN status messages. They facilitate load leveling and roaming within the premises LAN and allow sequence maintenance to improve security and performance within the NET. HELLO messages occur periodically in Access Intervals that contain priority SYNC messages. HELLOs are sent periodically relative to the sequence length, for instance, every 90 Access Intervals. HELLOs, like SYNC information, are optionally encrypted to provide greater security.

Each HELLO message includes a field for COST. COST is a measure of the base station to handle additional traffic. A device determining which of two or more base stations having adequate signal strength to register which will select the base with the lowest COST factor.

The base computes COST on the basis of how many devices are attached to the NET, the degree of bandwidth utilization, whether the base is wired or wireless, the number of frequencies experiencing consistent interference within the sequence, and the quality of the connection the base has within the premises LAN.

Fig. 14 illustrates the concept of base stations communicating neighboring base station information through HELLO messages to facilitate roaming of portable/mobile devices. In a premises LAN, base

- 71 -

stations 1401, 1403 and 1405 communicate SYNC information amongst themselves via wired backbone (LAN) 1407. In addition, a wireless base station 1409 (discussed above) similarly communicates with the base stations 1401, 1403 and 1405 via a wireless link 1411. A portable/mobile device 1413 is initially registered with base station 1401, which acts as a control point for the portable/mobile device 1413. HELLO messages transmitted by base station 1401 to portable/mobile device 1413 contain fields for neighboring base stations 1403, 1405 and 1409. These fields may indicate, for example, addresses of the neighboring bases, their COST, the hopping sequences, hopping sequence indices, number of Access Intervals per hop, and NET clock. The portable/mobile device 1413 detects the HELLOs transmitted from base station 1401 and uses the information for coarse synchronization with the other base stations 1403, 1405 and 1409. This permits the portable/mobile device to roam between base station coverage areas (i.e., between different NETs) without going through a full acquisition phase. Roaming of portable/mobile devices is discussed in more detail below.

Simply put, communication of neighbors' information permits each base station to advise its associated portable/mobile devices (i.e., those having common communication parameters) on how to capture HELLO

- 72 -

messages from neighboring base stations having different communication parameters. Such communication parameters may include, for example, hopping sequences, spreading codes, or channel frequencies.

5 For example, neighbors' information transmission is appropriate in any case where the system uses more than a single channel. For instance, in a direct sequence architecture, a single spreading code is often used. Capacity can be added to such a network by employing
10 different spreading codes at each base station. The neighbors' information included in the HELLO message from a given base station would include the spreading sequences of base stations providing coverage in adjacent coverage areas. Likewise, in a multiple frequency
15 channelized system, HELLO messages would include the channel frequencies of adjacent base stations.

 In addition to facilitating roaming, communication of neighbors' information may also facilitate the initial selection of a base station by a portable/mobile device
20 attaching to the premises LAN for the first time.

 Base station HELLO messages may also facilitate adaptive base station transmitter power control. For example, each base station HELLO transmission could specify the transmitter power level being used by the
25 base station. If a given attached portable/mobile device notes that the current base station transmitter power level is unnecessarily high (creating the possibility of

- 73 -

interference with other base stations), the portable/mobile unit could send a message to the base station indicating as such, and the base station could adjust the transmitter power level accordingly.

5 HELLO messages also enable communication of information indicating to all devices that certain changes in the NET are required. For example, the NET may switch hopping sequences periodically to improve security, or to avoid interference sources that
10 consistently interfere with one or two frequencies within a given sequence. Interference may result from outside sources, or from other NETs. Changes to the NET are communicated over the course of several HELLO messages (with a countdown) before the change occurs, so that all
15 devices are likely to be aware of changes and synchronize at the instant of change.

 In addition, if encryption is used, the encryption key may be periodically changed in HELLOs. Like hopping sequence changes, KEY changes are sent over several
20 HELLOs, and are encrypted using the existing key until the change goes into effect.

 As mentioned above, roaming portable and mobile computing devices operating in the premises LAN will routinely move between base station coverage areas. At
25 the maximum device velocity and expected coverage area per base station, a mobile device may be expected to cross a NET coverage contour in several seconds. Because

- 74 -

of the use of multiple, non-synchronized frequency hopping NETs, it is more difficult to provide for simple hand-off between base stations than it would be in a system that used cellular techniques with a single
5 frequency per cell. The premises LAN makes special provisions for roaming by transmitting coarse frequency hopping synchronization information in HELLO messages.

The premises LAN uses a spanning tree algorithm to maintain current information regarding the general
10 location of mobile devices within the network. When a device changes registration from one NET Control Point to another, routing information is updated throughout the infrastructure. Wired base stations may broadcast spanning tree updates to attached wireless base stations.

15 In the premises LAN, roaming portable and mobile devices initially select and register with a Base Station Control Point on the basis of link quality, i.e., signal quality, signal strength and COST information transmitted within HELLO messages. A device will remain attached to
20 a particular base station until the link quality degrades below an acceptable level, then it will attempt to determine if an alternative NET is available. Different device operating scenarios dictate different roaming strategies, discussed below.

25 An idle device monitors SYNC and HELLO messages from the Control Point device to maintain NET connectivity. Type 2 devices do not employ power management, and always

- 75 -

maintain their receivers in an active state. They monitor all SYNC messages. Type 1 and Type 3 devices typically employ power management, operating in standby or sleep modes of operation for many Access Intervals before activating their receivers for monitoring SYNC and HELLO messages. Control Points are guaranteed to send Priority SYNC frames every third Access Interval. HELLOs occur every 30th Priority SYNC frame. Power managed devices employ sleep algorithms synchronized to wake for the minimum period necessary to guarantee receipt of priority SYNC, HELLO, and Pending Message transmissions before resuming SLEEP.

Type 2 devices are typically operated from high capacity vehicular power systems, which eliminates the need for power management. These devices may travel at velocities near the maximum system design specification, dictating more frequent roaming. Type 2 devices will initiate a search for an alternative NET if SYNC messages are consistently received at signal strengths below a Roaming Threshold or if reception errors are consistently detected. Because of the effects of frequency selective fading, signal strength information is averaged over the course of several hops within the hopping sequence.

If roaming is indicated, the device initiates a Roaming Algorithm, using Neighbors' information from the most recent HELLO to attempt synchronization with another candidate NET. If SYNC is not detected within 6 hops,

- 76 -

another candidate from the Neighbors list will be selected, and the process repeated. Once SYNC is attained on an alternative NET, the device will monitor signal strength and data errors for several hops to determine link quality. If link quality is acceptable, the device will continue monitoring until a HELLO is received. If COST is acceptable, it will then register with the new NET. The Control Point device will update the Spanning Tree over the wired backbone (or by RF if a wireless base). If link quality or COST is unacceptable, another candidate from the Neighbors list is selected and the process repeated. This continues until an acceptable connection is established. If a connection cannot be established, the device must return to the original NET or employ the initial acquisition algorithm.

Type 2 devices also have the option of monitoring other NETs before degradation of their NET connection. They may do so by monitoring their own NET for the SYNC and pending message list transmissions, then scanning other candidate NETs during the Sessions period of their NET. Other type devices may do so less frequently.

Type 1 and Type 3 devices may sleep extensively when idle, preferably activating every nine Access Intervals to resynchronize and check pending messages. Successful reception of at least one SYNC during three monitoring periods is necessary to maintain fine synchronization to the NET clock. Failure to receive two of three SYNC

- 77 -

frames, or receipt of two or three SYNC messages with poor signal strength are possible indications of the need to further test link quality by remaining active for several consecutive SYNC transmissions. If signal strength or data errors over several hops indicates that link quality is poor, or if a received HELLO message indicates high COST, the roaming algorithm is initiated, and alternative NETs are evaluated, as in the case of Type 2 devices.

Some battery powered devices may sleep for periods of time more than nine Access Intervals. For example, devices with extremely limited battery capacity may sleep between HELLOs, or several HELLO periods, after which they must remain active for several consecutive Access Intervals to regain fine synchronization and assess whether to initiate roaming.

A Type 1, Type 2, or Type 3 device that has inbound message requirements immediately activates its receiver and waits for a SYNC and subsequent Reservation Opportunities. A device that does not detect SYNC messages over the course of six Access Intervals immediately initiates the Roaming Algorithm.

Outbound messages for devices that have changed coverage areas, but which have not yet registered with a new Control Point device, are problematic. For example, in the premises LAN, messages will be forwarded to the Base Station that the device had previously been attached

- 78 -

to. The base station may attempt to poll the device during one or more Access Intervals, then transmit the unit address in the pending message list periodically for several seconds before disregarding it. Once the unit
5 attaches to a base, the message must be transferred from the previous base station for delivery to the unit. All of these activities require transmission bandwidth on either the backbone or RF media, waste processing resources within the premise LAN, and result in delayed
10 delivery.

As this premises LAN embodiment is designed, the network has no means of distinguishing messages it cannot deliver due to roaming from messages that should be
15 retried due to signal propagation characteristics, interference, or sleeping devices. For this reason, the roaming algorithm may be designed to allow devices to quickly detect that they have lost connectivity within their current NET, and re-attach to a more favorably located base station.

20 Some improvement in delivering pending messages to roaming terminals can be obtained by routinely propagating pending message lists over the wired backbone. When a device attaches to a base station, that base is able to immediately ascertain that the device has
25 a pending message, and initiate forwarding of the message for delivery to the device.

- 79 -

In the preferred frequency hopping embodiment of the present invention, the hopping sequence consists of $3m + 1$ frequencies, where m is an integer. 79 frequencies are preferred. This embodiment will support hopping rates of 100, 50 hops per second at 1 Access Interval per dwell, 25 hops per second at 2 frames per dwell, and 12.5 hops per second at 4 frames per dwell. Other rates can be supported for other Access Interval Durations. For example, if the Access Interval is optimized to 25 ms, hop rates of 80, 40, 20, and 10 hops per second would be supported.

All devices within the NET may have one or more hopping tables that contain potential hopping sequences that may be used. Up to 64 sequences may be stored in each device. Each sequence has an identifier, and each frequency in each sequence has an index. The sequence identifier and index are communicated in the SYNC transmission.

All SYNC transmissions may be block encrypted to prevent unauthorized devices from readily acquiring hopping synchronization information. To facilitate encryption, the encryption key may initially be factory set to a universal value in all devices. Users would then have the option of changing this key, by providing a new key to each device in the system. This may be accomplished through keyboard entry or other secure means. Keys may also be changed through the NET.

- 80 -

To facilitate hopping management, a hopping control portion of a protocol controller will download a hopping table to a radio modem, and will signal the radio modem when to hop. This approach consolidates timing functions in the protocol controller, while not requiring the controller to be concerned with conveying frequency selection data to the modem each hop.

The NET may switch hopping sequences periodically to improve security, or to avoid interference sources that consistently interfere with one or two frequencies within a given sequence. As mentioned above, changes to the NET are communicated over the course of several HELLO messages before the change occurs so that all devices are likely to be aware of changes.

Initial synchronization requires devices to ascertain the hopping sequence, the hop rate, and the specific frequency from the hopping sequence currently in use. Synchronization information is contained in two types of routine messages. The SYNC field at the beginning of an Access Interval contains synchronization information including the hopping sequence, the index of the current frequency within the sequence, the number of Access Intervals per hop, and the length of the Access Interval. It also contains a timing character that communicates the NET master clock to all listening devices. Termination messages in the Sessions period,

- 81 -

ACK and CLEAR, contain the same information, but do not contain the timing character.

The simplest method for attaining synchronization is to *Camp* -- select a quiet frequency that is likely to be within a sequence in use -- and listen for valid synchronization information. If a SYNC message is detected, the listening device immediately has both coarse and fine synchronization, and can begin the registration process.

If SYNC is not detected, but a termination message is, then the device has acquired coarse synchronization. The particulars of the hopping sequence are known, but the boundaries of the dwells are not. To acquire fine synchronization, it begins hopping at the indicated hopping rate, listening for SYNC. If SYNC is not detected after a reasonable number of hops, preferably 12 or 15, the device reverts to camping.

The worst case scenario for synchronization is to synchronize to a single NET that is idle. Given a 79 frequency hopping sequence, one Access Interval per hop, and SYNC transmissions every third Access Interval if the NET is idle, it may take nine cycle times to guarantee that a SYNC transmission will be detected with 99.5% probability. At 50 hops per second, synchronization could require as long as 14 seconds. At 100 hops per second, 7 seconds is required.

- 82 -

At 2 Access Intervals per hop, a SYNC transmission is guaranteed to occur every frequency over 2 cycles of the hopping sequence. Six cycles are required for 99.5% probability of acquisition, corresponding to 19 seconds at 25 hops per second.

At 4 Access Intervals per hop, at least one SYNC is guaranteed to occur each hop. Three cycles of the hopping sequence are required for 99.5% acquisition probability. At 12.5 hops per second, this also requires 19 seconds.

This illustrates the advantage of scalability. A device that uses an acquisition algorithm suitable for 2 or 4 Access Intervals per hop will also acquire a NET that hops at 1 Access Interval per hop. The algorithm may be as follows:

1. The device scans candidate frequencies until it finds one with no Received Signal Strength Indicator indication.
2. The device remains on the frequency for 6.32 seconds 2 Access Interval/hop @ 25 Hops/second x 2, or 4 Access Interval/hop @ 12.5 hops/second x 1, or until it detects a SYNC message or a valid termination message.
3. If SYNC is detected, the device synchronizes its internal clock to the SYNC, and begins hopping with the NET for the next 11 hops. It may attempt registration after detecting valid SYNC and any Reservation Opportunity. If synchronization is not verified by detection of SYNC within the 11 hops, the acquisition algorithm is reinitialized.
4. If a message termination (either an ACK or CLEAR) is detected, the device immediately hops to the next frequency in the sequence and waits for the SYNC. It is coarsely synchronized to

- 83 -

the NET but has a timing offset from the NET clock.

When the next SYNC is received, the device synchronizes its clock to the NET clock and initiates registration. If SYNC is not received within a dwell time, the device hops to the next frequency in sequence. This continues until SYNC is attained, or until 15 hops have passed without receiving SYNC, after which the acquisition sequence is restarted.

- 5 5. If coarse acquisition is not obtained within 6.3 seconds, the device selects another frequency and repeats the process beginning with step 2.

Camping provides a worst case acquisition performance that is perceptibly slow to the human user of a portable device. The preferred approach has the receiver scan all potential frequencies in ascending order, at 125 μ sec increments. When the highest frequency is reached, the search begins again at the lowest frequency. The 125 μ s sampling rate is much faster than the 250 μ sec channel switching time specification of the RF modem. This is possible because the overall switching time specification applies to worst case frequency switching intervals, i.e., from the highest to the lowest operating frequency. By switching a single channel at a time, switching may be maintained over frequency intervals very near a synthesizer phase detectors' phase lock range, allowing nearly

- 84 -

instantaneous frequency switching. The change from highest to lowest frequency at the end of the scan requires the standard 250 μ sec.

5 The 125 μ sec monitoring interval allows 85 μ s to ascertain if receive clock has been detected prior to switching to the next frequency. The monitoring interval should be selected to be non-periodic with respect to the access interval. For example, the 125 μ sec interval allows the entire hopping sequence to be scanned 2(n + 1)
10 times in a 20 ms access interval.

If clock is recovered at any frequency, the receiver remains on frequency for a Reservation Opportunity and initiates channel access through the procedure described above. The scanning approach is less deterministic in
15 terms of acquisition probability than camping, but the search time required for 99.5% acquisition probability is about 80 Access Intervals, or three times faster than that for camping.

A hybrid approach that scans only three or four
20 consecutive frequencies incorporates the deterministic aspects of camping with some of the improved performance of the scanning algorithm. For scanning over a small number of frequencies an up/down scan is preferred, i.e., 1,2,3,2,1,2,3 since all frequency changes can be
25 accomplished at the faster switching rate. The end frequencies are visited less often than those in the center. The number of frequencies used, e.g., 3 or 4, is

- 85 -

selected so that all can be scanned during the preamble duration of a minimum length transmission.

All devices are required to have unique 48 bit global addresses. Local 16 bit addresses will be assigned for reduced overhead in communications. Local addresses will not be assigned to devices whose global addresses are not on an authentication list maintained in each base station and routinely updated over the infrastructure.

Once a device has attained synchronization, it must register with the control point to be connected with the NET. It initiates this by sending a Request for Poll indicating a registration request, and including its global address. The control point will register the device, and provide a short. Network Address as an outbound message. The Control point will generate the short address if it is a single NET, or exchange the global address for a short Network Address with a Network Address Server if the NET is part of a larger infrastructured network of a premises LAN.

Once a device is synchronized to a NET, it must periodically update its local clock to the NET clock communicated in the SYNC message. The SYNC message contains a character designated as the SYNC character that transfers the NET clock synchronization. This may be the beginning or ending FLAG in the SYNC message, or a specific character within the message.

- 86 -

The maximum expected frequency error between NET and device local clocks is 100 parts per million. To maintain a 50 μ s maximum clock error, the local device clock must be re-synchronized at 500 ms intervals. At 20 ms per access interval, a non-sleeping device has up to 26 SYNC opportunities within that period in which to re-synchronize and maintain required accuracy.

As mentioned above, it is desirable that battery powered devices have the capability to sleep, or power off, for extended periods of time to conserve power. The term sleeping terminal in this instance may refer to a device that powers down its radio communication hardware to save power while maintaining other functions in an operational state, or a device that power manages those functions as well. In the power managed state, the device must maintain its hop clock so that full acquisition is not required every time power management is invoked.

Devices that must sleep to manage their power consumption use *Priority SYNC Messages* to maintain synchronization. Priority SYNC Messages occur every three Access Intervals. In times of low NET activity, non-priority SYNC messages are omitted. By coordinating power management with Priority SYNC Messages, power managed devices can be guaranteed to wake up for Access Intervals where SYNCs will be present, even if the NET activity is low during the sleep period.

- 87 -

A sleeping device with no transmission requirements may sleep for eight 20 ms access intervals, and wake only for the SYNC and Reservation Poll at the beginning of the ninth Access Interval to monitor pending messages before
5 returning to the sleep state, for a duty cycle of less than 5%. This provides three opportunities to synchronize to the NET clock within a 540 ms window. A flow chart depicting the a device sleeping for several access intervals is shown in Fig. 17.

10 Devices may also sleep for longer periods of time, at the risk of losing fine synchronization. They may compensate by advancing their local clocks to account for the maximum timing uncertainty. For example, a terminal could sleep for 5 seconds without re-synchronizing by
15 waking up 500 microseconds before it expects an Access Interval to begin, and successfully receive SYNC messages. This technique is valid for extended periods of time, up to the point where the maximum timing error approaches 50% of an Access Interval. A flow chart
20 depicting the a device sleeping for several seconds is shown in Fig. 18.

A power managed device that requires communication during a sleep period may immediately wake and attempt access to the NET at the next available Reservation
25 Opportunity.

A device requiring communications may be able to register with one of several NETs operating in its

- 88 -

vicinity, with transmissions occurring on many frequencies simultaneously. A good strategy is to synchronize to a NET that provides an acceptable communication link, then monitor HELLO messages to
5 determine other candidate NETs before attaching to a particular NET by registering with the control point device.

As described above, a spontaneous wireless local area network or spontaneous LAN is one that is
10 established for a limited time for a specific purpose, and which does not use the premises LAN to facilitate communications between devices or provide access to outside resources. Use of spontaneous LAN allows portable devices to share information, files, data, etc.,
15 in environments where communication via the premises LAN is not economically justifiable or physically possible. A spontaneous LAN capability also allows portable/mobile devices to have an equally portable network. Peripheral and vehicular LANs are examples of such spontaneous LANs.

20 Requirements for spontaneous LAN differ from an infrastructured premises LAN in several significant areas. The number of devices in a spontaneous LAN is likely to be smaller than the number that a single NET in a premises LAN must be capable of supporting. In
25 addition, coverage areas for spontaneous LANs are typically smaller than coverage areas for a base station participating in the premises LAN. In a spontaneous LAN,

- 89 -

communication often takes place over relatively short distances, where devices are within line of sight of each other.

5 In an premises LAN, the majority of communications are likely to involve accessing communication network resources. For example, portable devices with limited processing capabilities, memory, and power supplies are able to access large databases or powerful computing engines connected to the AC power grid. Base stations
10 within the premises LAN are well suited to the role of Control Points for managing synchronization and media access within each NET.

In a spontaneous LAN, however, communications are limited to exchanges with spontaneous NET constituents.
15 Additionally, NET constituents may potentially leave at any time, making it difficult to assign control point responsibilities to a single device. A shared mechanism for synchronization and media access is preferable in most cases.

20 In a spontaneous LAN, battery power limitations may preclude assignment of a single device as a control point. The routine transmission of SYNC and access control messages places a significant power drain on a portable, battery powered device. Also, the control
25 point architecture dictates that transmissions intended for devices other than the control point be stored and

- 90 -

forwarded to the destination device, further increasing battery drain, and reducing system throughput.

Moreover, the use of scheduled transmission in a premises LAN is likely to differ from use in a spontaneous LAN. For example, unlike the premises LAN, in the spontaneous LAN, applications such as massaging and two way voice communications may only occasionally be used, whereas video transmission and telemetry exchange may be prevalent.

To promote compatibility and integration with the premises LAN, operational differences required by multiple participating devices should be minimized. For example, selecting relatively close frequency bands for each LAN aids in the design of a multiple LAN transceiver, reducing circuitry, cost, power, weight and size while increasing reliability. Similarly, selecting communication protocols so that the spontaneous LAN protocol constitutes a subset or superset of premises LAN may enable a given device to more effectively communication in both LANs, while minimizing both the overall protocol complexity and potentially limited memory and processing power.

Use of frequency hopping is desirable in premises LAN because of its ability to mitigate the effects of interference and frequency selective fading. In the case of the latter, frequency hopping allows systems to be installed with less fade margin than single frequency

- 91 -

systems with otherwise identical radio modem characteristics, providing improved coverage.

The potentially smaller coverage area requirement of spontaneous LANs, however, allows single frequency operation to be considered for some applications, e.g., such as a peripheral LAN. Regulatory structures are in place in some countries to allow single frequency operation in the same bands as frequency hopping systems, providing that single frequency devices operate at reduced power levels. The lower transmit power of single frequency operation and elimination of periodic channel switching are desirable methods of reducing battery drain. The choice of single frequency or frequency hopped operation is dictated by the coverage requirements of the network, and may be left as an option to device users.

As noted earlier, the basic Access Interval structure is suited to single frequency operation as well as to frequency hopping. SYNC messages in a single frequency system substitute a single frequency indication in the hopping sequence identifier field.

A spontaneous LAN comes into existence when two or more devices establish communications, and ceases when its population falls to less than two. Before a spontaneous LAN can be established, at least two devices must agree upon a set of operating parameters for the network. Such agreement may be pre-programmed else

- 92 -

exchanged and acknowledged prior to establishing the spontaneous LAN. Once the spontaneous LAN is established, other devices coming into the network must be able to obtain the operating parameters and acquire access.

5 More specifically, to establish a spontaneous LAN, a computing device must first identify at least one other network device with which spontaneous LAN communication is desired. To identify another network device, the
10 computing device may play an active or passive role. In an active role, the computing device periodically broadcasts a request to form spontaneous LAN with either a specific network device or, more likely, with a specific type of network device. If a network device
15 fitting the description of the request happens to be in range or happens into range and is available, it responds to the periodic requests to bind with the computing device, establishing the spontaneous LAN. Alternately, the network device may take a passive role in
20 establishing the spontaneous LAN. In a passive role, the computing device merely listens for a request to form a spontaneous LAN transmitted by the appropriate network device. Once such a network device comes into range, the computing device responds to bind with the network
25 device, establishing the spontaneous LAN.

The choice of whether a device should take a passive or active role is a matter of design choice. For

- 93 -

example, in one embodiment where peripheral devices have access to AC power, the roaming computer terminals take a passive role, while the peripheral devices take a more active role. Similarly, in another embodiment where a vehicle terminal has access to a relatively larger battery source, an active role is taken when attempting to form a spontaneous LAN, i.e., a vehicular LAN, with a hand-held computing device.

Binding, a process carried out pursuant to a binding protocol stored in each network device, may be a very simple process such as might exist when creating a spontaneous LANs that operates on a single frequency channel. Under such a scenario, a simple acknowledge handshake between the computing terminal and the other network device may be sufficient to establish a spontaneous LAN pursuant to commonly stored, pre-programmed operating parameters. However, more complex binding schemes may also be implemented so as to support correspondingly more complex spontaneous LANs as proves necessary. An example of a more complex binding scheme is described below.

It is desirable in some large spontaneous LANs for one device to be designated as a fully functional control point, providing identical NET operation to a single NET in the premises LAN. Providing that all devices share a hopping table and encryption key, the designated device would initiate control point activities, and other

- 94 -

devices would synchronize to the designated unit. A device with greater battery capacity, or one that can be temporarily connected to AC power is best suited to the dedicated control point function. This architecture is applicable to *Client-Server* applications (where the server assumes the control point function), or to other applications where a single device is the predominant source or destination of communications. A portable device used as a dedicated control point is required to have additional programming and memory capacity to manage reservation based media access, pending message lists, and scheduled service slot allocations.

In embodiments where communication requirements of a spontaneous LAN are largely peer to peer, there may be no overwhelming candidate for a dedicated Control Point. Thus, in such cases, the Control Point function is either distributed among some or all the devices within the spontaneous LAN. In such scenarios, the interleaved Access Interval approach used for wireless base stations is employed. Initially, control point responsibilities are determined during the binding process. Users may designate or redesignate a Control Point device when several candidates are available.

For spontaneous LANs, access intervals may be simplified to reduce power consumption, program storage and processing power requirements for portable devices used as control points. Control Point devices transmit

- 95 -

SYNC, pending message lists, and Time Division Multiple Access slot reservations normally, but only use the single slot reservation Poll (Idle Sense Multiple Access). The reservation poll contains a field
5 indicating reduced control point functionality. This places other devices in a point-to-point communication mode, using the Implicit Idle Sense Algorithm. The probability factor p communicated in the reservation poll is used for the Implicit Idle Sense algorithm. Control
10 point devices may use the deferred SYNC mechanism for light system loading, transmitting Priority SYNC every third Access Interval to further decrease their transmission requirements. Control point devices must monitor the reservation slot for messages addressed to
15 them, but may sleep afterwards.

Request for Polls initiated under Implicit Idle Sense use point-to-point addressing, indicating the address of the destination device directly, rather than the control point device. This eliminates the need for
20 the Control Point device to store and forward transmissions within the spontaneous LAN. The device detecting its address in a Request for Poll begins a session, after employing the Implicit Idle Sense algorithm, by Polling the source address identified in
25 the Request for Poll. The terminating ACK and CLEAR messages contain an Explicit Idle Sense probability factor equal to that in the original reservation poll.

- 96 -

To allow for power managed devices, the Control Point device maintains a pending message list. Devices that have been unable to establish communication with a sleeping device initiate a session with the Control Point device to register the pending message. Upon becoming active, the sleeping device will initiate a Poll to the device originating the pending message. The Control Point device will eliminate the pending message indication by aging, or by receipt of communication from the destination device clearing the pending message. Control point devices are not required to store pending messages, only addresses.

As mentioned above, HELLO messages are broadcast to indicate changes in NET parameters. HELLO messages may be omitted to simplify the Control Point function in spontaneous LANs.

Devices are assigned local addresses upon registration with the Control Point device. Devices may communicate an alias that identifies the device user to other users to the Control Point device where it is stored in an address table. The address table may be obtained by other network constituents by querying the Control Point device. A peripheral LAN is a type of spontaneous LAN which serves as a short range interconnect between a portable or mobile computing device (MCD) and peripheral devices.

- 97 -

Designers of portable products are constantly challenged with reducing size, weight, and power consumption of these devices, while at the same time increasing their functionality and improving user ergonomics. Functions that may be used infrequently, or which are too large to fit within the constraints of good ergonomic design may be provided in peripheral devices, including printers, measurement and data acquisition units, optical scanners, etc. When cabled or otherwise physically connected to a portable product, these peripherals often encumber the user, preventing freedom of movement or mobility. This becomes more problematic when use of more than one peripheral is required.

A second consideration for portable product design is communication docking. A communication dock is a device that holsters or houses a portable unit, and provides for communication interconnection for such tasks as program downloading, data uploading, or communication with large printers, such as those used for printing full sized invoices in vehicular applications. Communication docking of a portable unit may also involve power supply sharing and/or charging.

The requirement for communication docking capability forces newer portable product designs to be mechanically compatible with older docking schemes, or may require that new docks, or adapters, be developed for each new generation of portable device. Product specific docking

- 98 -

approaches eliminate compatibility between devices manufactured by different suppliers. This has hindered development of uniform standards for Electronic Data Interchange between portable devices and fixed computing systems.

Physical connection between a portable device with a peripheral or communication dock also hinders user efficiency. Peripheral devices are generally attached with cable. If a peripheral is small enough to be carried or worn on a belt, the mobility of the user may be maintained. If a user must carry a hand-held portable device that is connected to a belt mounted peripheral the assembly cannot be set down while a task that requires movement to a location several feet away is undertaken unless the portable device and peripheral are disconnected. Likewise, connection to peripherals too large to be portable requires the user to frequently connect and disconnect the device and the peripheral.

Use of wireless peripheral LAN interconnection greatly simplifies the task of portable devices communicating with peripherals. In doing so, wireless connectivity allows improved ergonomics in portable product design, flexibility in interconnection to one or more peripherals, freedom of movement over a radius of operation, forward and backward compatibility between portable units and peripherals, and potential

- 99 -

communications among products manufactured by different vendors.

Constituents within a peripheral LAN generally number six or fewer devices. One roaming computing device and one or two peripherals comprise a typical configuration. Operating range is typically less than fifty feet.

Because the computing devices generally control the operation of peripheral devices, in a peripheral LAN a master/slave type protocol is appropriate. Moreover, roaming computing devices serving as master are well suited to the role of Control Points for managing synchronization and media access within each peripheral LAN. All peripheral communications are slaved to the master.

In a peripheral LAN, roaming mobile or portable computing devices and wireless peripherals may all operate from battery power. Operating cycles between charging dictate use of power management techniques.

Although all participants in a peripheral LAN might also be configured to directly participate in the premises LAN, the trade-offs in cost, power usage and added complexity often times weighs against such configuration. Even so, participants within a peripheral LAN can be expected to function in a hierarchical manner, through a multiple participating device, with the premises LAN. Thus, the use of a much simpler, lower-

- 100 -

power transceiver and associated protocol may be used in the peripheral LAN.

As previously described, a roaming computing device serving as a master device may itself be simultaneously attempting to participate in other networks such as the premises or vehicular LANs. Considerable benefits arise if the radio and processing hardware that supports operation within the wireless network can also support such operation. For example, a device that is capable of frequency hopping is inherently suited to single frequency operation. If it can adjust transmitter power level and data rate to be compatible with the requirements of the peripherals LAN, it can function in both systems. The major benefits of common transceiver hardware across LANs include smaller product size, improved ergonomics, and lower cost.

Specifically, in one embodiment, radio communication on the premises LAN, as described herein, takes place using radio transceivers capable of performing frequency-hopping. To communicate on a peripheral LAN, such transceivers could also utilize frequency-hopping at a lower power. However, such transceivers are relatively expensive in comparison to a lower power, narrow-band, single frequency transceivers. Because of the cost differential, it proves desirable to use the single frequency transceivers for all peripheral devices which will not participate in the premises LAN. Therefore, the

- 101 -

more expensive, frequency-hopping transceivers which are fitted into roaming computing devices are further designed to stop hopping and lock into the frequency of the single frequency transceiver, allowing the
5 establishment of peripheral LANs.

Instead of frequency hopping, the peripheral LAN may also use narrow-band, single frequency communication, further simplifying the radio transceiver design for commonality. In another embodiment of the peripheral LAN
10 transceivers, operation using one of a plurality of single frequency channels is provided. Thus, to overcome interference on one channel, the transceiver might select from the remaining of the plurality an alternate, single operating frequency with lesser channel interference. To
15 accommodate the plurality of single frequency channels, the peripheral LAN transceivers may either communicate an upcoming frequency change so that corresponding peripheral LAN participants can also change frequency, or the transceivers may be configured to use frequency
20 synthesis techniques to determine which of the plurality a current transmission happens to be.

The Access Interval structure is also an appropriate choice for peripheral LAN operations. In one embodiment, to provide for simplicity and tighter integration, the
25 Access Interval for the peripheral LAN is a subset of the Access Interval used in the premises LAN. HELLO messages, Implicit Idle Sense, Data Rate Switching, and

- 102 -

scheduled services are not implemented. Peripheral devices normally sleep, activate their receivers for SYNC transmissions from the participating master device, and resume sleeping if no pending messages are indicated and they have no inbound transmission requirements. Access Intervals occur at regular intervals, allowing for power management. Access Intervals may be skipped if the master has other priority tasks to complete.

To initialize the peripheral LAN, a device desiring initialization, a master device, selects a single operating frequency by scanning the available frequencies for one with no activity. A typical master device might be a roaming computing device desiring access to a local peripheral. Default values for other parameters, including Access Interval duration, are contained within each participant's memory. Such parameters may be pre-adjusted in each participant to yield specific performance characteristics in the peripheral LAN.

Once a master device identifies a single frequency, slaves, which are generally peripherals, are brought into the peripheral LAN through a process called binding. Binding is initiated by the master device by invoking a binding program contained therein. Slaves, such as peripherals, are generally programmed to enter a receptive state when idle. Thus, in one embodiment, the master device accomplishes binding by transmitting Access Intervals of known duration sequentially on a series of

- 103 -

four frequencies spread throughout the available frequency range. The specific frequencies and Access Interval durations used are stored as parameters in all potential participating devices. A 250 Kbps transfer rate is appropriate in some embodiments of the peripheral LAN, reflecting a balance between performance and complexity in peripheral devices.

A slave, e.g., a peripheral, responds to the binding attempts by the master device on a given frequency until the slave successfully receives and establishes communication with the master device. If they do not establish communication after four Access Intervals, the slave switches to the next frequency for four Access Interval periods. Once communication is established, the slave registers with the master and obtains the master device's selected operating frequency and related communication parameters. When all slave devices have been bound, the master terminates the binding program and normal operation at the selected single frequency may begin.

Referring to Fig. 15, in a hierarchical network, peripheral LAN masters use a secondary access interval 1501 that is synchronized to the Access Interval of a parent (premises) LAN control point. Peripheral LAN Access Intervals occur less frequently than premises LAN Access Intervals, e.g., every other or every third Priority SYNC Access Interval.

- 104 -

During the premises LAN Access Interval, the peripheral LAN master device monitors the premises LAN control point for SYNC 1503 reservation poll 1505 and exchanges inbound and outbound message according to the normal rules of the access protocol. The master switches to the peripheral LAN frequency, and transmits its own SYNC frame 1507 during the session period 1509 of its parent control point allowing communication with its peripherals. The peripheral LAN Access Interval is generally shorter than the premises LAN Access Interval, so that it does not extend beyond the premises LAN Access Interval boundary. At the end of the peripheral LAN Access Interval 1501, the master switches to the premises LAN frequency for the next SYNC 1503.

The secondary SYNC 1507 may only be transmitted if the peripheral LAN master is not busy communicating through the premises LAN. If a communication session is occurring, the master must defer SYNC, preventing communication with its peripherals during that Access Interval. The master must also defer SYNC if the current frequency in the LAN is prone to interference from the peripheral LAN frequency, i.e., they are the same frequency or adjacent frequencies. If two consecutive SYNCs are deferred, peripherals will activate their receivers continuously for a period of time, allowing the master to transmit during any Access Interval. This approach is also applicable when the master roams between

- 105 -

frequency hopping NETs. Since NETs are not synchronized to one another, the devices in the peripheral LAN adjust Access Interval boundaries each time the master roams. If peripherals do not detect SYNC within a time-out
5 period, they may duty cycle their reception to conserve battery power.

Referring to Figs. 16a and 16b, a Roaming Algorithm Flow Diagram illustrates how a roaming computing device will select a suitable base station. Roaming computing
10 devices operating in the infrastructured network environment formed by the base stations will routinely move between base station coverage areas. The roaming computing devices are able to disconnect from their current base station communication link and reconnect a
15 communication link to a different base station, as necessitated by device roaming.

Base stations transmit HELLO messages to devices in their coverage area. These HELLO messages communicate to roaming computing devices the cost of connection through
20 the base station, addresses of neighboring base stations, and the cost of connection through these neighboring base stations. This information allows roaming computing devices to determine the lowest cost connection available and to connect to the base station with the lowest cost.

25 In addition, base station HELLO message may include communication parameters of neighboring base stations, such as frequency hopping sequences and indices, spread

- 106 -

spectrum spreading codes, or FM carrier channel frequencies. This information allows roaming computing devices to roam and change base station connections without going through a full acquisition phase of the new base station's parameters.

5 Roaming computing devices initially select and register with a base station control point on the basis of link quality: signal strength and cost information transmitted within HELLO messages. A device will remain
10 attached to a particular base station until the link quality degrades below an acceptable level; then it will attempt to determine if an alternative base station connection is available. The device initiates a roaming algorithm, using neighbors information from the most
15 recent HELLO message to attempt connection with another candidate base station. If connection fails, another candidate from the neighbors list will be selected, and the process repeated. Once connection is made with an alternative base station, the device will monitor signal
20 strength and data errors to determine link quality. If link quality is acceptable, the device will continue monitoring until a HELLO message is received. If the cost is acceptable, it will register with the new base station, and the base station will update the spanning
25 tree over the infrastructure. If link quality or cost is unacceptable, another candidate from the neighbors list is selected and the process repeated. This continues

- 107 -

until an acceptable connection is established. If one cannot be established, the device must return to the original base station connection or employ the initial acquisition algorithm.

5 Fig. 28a illustrates an embodiment of the hierarchical communication system according to the present invention communication is maintained in a warehouse environment. Specifically, a worker utilizes a roaming computing device, a computer terminal 3007, and
10 a code reader 3009 to collect data such as identifying numbers or codes on warehoused goods, such as the box 3010. As the numbers and codes are collected, they are forwarded through the network to a host computer 3011 for storage and cross-referencing. In addition, the host
15 computer 3011 may, for example, forward cross-referenced information relating to the collected numbers or codes back through the network for display on the terminal 3007 or for printing on a printer 3013. Similarly, the collected information may be printed from the computer
20 terminal 3007 directly on the printer 3013. Other exemplary communication pathways supported include message exchanges between the computer terminal 3007 and other computer terminals (not shown) or the host computer 3011.

25 The host computer 3011 provides the terminal 3007 with remote database storage, access and processing. However, the terminal 3007 also provides for local

- 108 -

processing within its architecture to minimize the need to access the remote host computer 3011. For example, the terminal 3007 may store a local database for local processing. Similarly, the terminal 3007 may run a variety of application programs which never, occasionally or often need access to the remote host computer 3011.

Many of the devices found in the illustrative network are battery powered and therefore must conservatively utilize their radio transceivers. For example, the hand-held computer terminal 3007 receives its power from either an enclosed battery or a forklift battery (not shown) via a communication dock within the forklift 3014. Similarly, the code reader 3009 operates on portable battery power as may the printer 3013. The arrangement of the communication network, communication protocols used, and data rate and power level adjustments help to optimize battery conservation without substantially degrading network performance.

In the illustrated embodiment shown in Fig. 28a, the hierarchical communication system of the present invention consists of a premises LAN covering a building or group of buildings. The premises LAN in the illustrated embodiment includes a hard-wired backbone LAN 3019 and base stations 3015 and 3017. A host computer 3011 and any other non-mobile network device located in the vicinity of the backbone LAN 3019 can be directly attached to the backbone LAN 3019. However, mobile

- 109 -

devices and remotely located devices must maintain connectivity to the backbone LAN 3019 through either a single base station such as the base station 3015, or through a multi-hop network of base stations such as is illustrated by the base stations 3015 and 3017. The base stations 3015 and 3017 contain a relatively higher power transmitter, and provide coverage over the entire warehouse floor. Although a single base station may be sufficient, if the warehouse is too large or contains interfering physical barriers, the multi-hop plurality of base stations 3017 may be desirable. Otherwise, the backbone LAN 3019 must be extended to connect all of the base stations 3017 directly to provide sufficient radio coverage. Through the premises LAN, relatively stable, longer range wireless and hard-wired communication is maintained.

Because roaming computing devices, such as the handheld computer terminal 3007, cannot be directly hard-wired to the backbone LAN 3019, they are fitted with RF transceivers. To guarantee that such a network device can directly communicate on the premises LAN with at least one of the base stations 3015 and 3017, the fitted transceiver is selected to yield approximately the same transmission power as do the base stations 3015 and 3017. However, not all roaming network devices require a direct RF link to the base stations 3015 and 3017, and some may not require any link at all. Instead, with such devices,

- 110 -

communication exchange is generally localized to a small area and, as such, only requires the use of relatively lower power, short range transceivers. The devices which participate in such localized, shorter range communication form spontaneous LANs.

For example, the desire by a roaming terminal to access peripheral devices such as the printer 3013 and modem 3023, results in the roaming terminal establishing a peripheral LAN with the peripheral devices. Similarly, a peripheral LAN might be established when needed to maintain local communication between a code scanner 3009 and the terminal 3007. In an exemplary embodiment, the printer 3013 are located in a warehouse dock with the sole assignment of printing out forms based on the code information gathered from boxes delivered to the dock. In particular, as soon as the code reader gathers information, it relays the information along a peripheral LAN to the terminal 3007. Upon receipt, the terminal 3007 communicates via the premises LAN to the host computer 3011 to gather related information regarding a given box. Upon receipt of the related information, the terminal 3007 determines that printing is desired with the printer 3013 located at the dock. When the forklift 3014 enters the vicinity of the dock, the terminal 3007 establishes a peripheral LAN with the printer 3013 which begins printing the collected code information.

- 111 -

To carry out the previous communication exchange, the printer 3013 and code reader 3009 are fitted with a lower power peripheral LAN transceivers for short range communication. The computer terminal 3007 transceiver is not only capable of peripheral LAN communication, but also with the capability of maintaining premises LAN communication. In an alternate exchange however, the code reader 3009 might be configured to participate on both LANs, so that the code reader 3009 participates in the premises LAN to request associated code information from the host computer 3011. In such a configuration, either the code reader 3009 or terminal 3007 could act as the control point of the peripheral LAN. Alternately, both could share the task.

With capability to participate in the peripheral LAN only, the code reader 3009, or any other peripheral LAN participant, might still gain access to the premises LAN indirectly through the terminal 3007 acting as a relaying device. For example, to reach the host computer 3011, the code reader 3009 first transmits to the computer terminal 3007 via the peripheral LAN. Upon receipt, the computer terminal 3007 relays the transmission to one of the base stations 3015 and 3017 for forwarding to the host 3011. Communication from the host 3011 to the code reader 3009 is accomplished via the same pathway.

It is also possible for any two devices with no access to the premises LAN to communicate to each other.

- 112 -

For example, the modem 3023 could receive data and directly transmit it for printing to the printer 3013 via a peripheral LAN established between the two. Similarly, the code reader 3009 might choose to directly communicate
5 code signals through a peripheral LAN to other network devices via the modem 3023.

In an alternate configuration, a peripheral LAN base station 3021 is provided which may be directly connected to the backbone LAN 3019 (as shown) or indirectly
10 connected via the base stations 3015 and 3017. The peripheral LAN base station 3021 is positioned in the vicinity of other peripheral LAN devices and thereafter becomes a control point participant. Thus, peripheral LAN communication flowing to or from the premises LAN
15 avoids high power radio transmissions altogether. However, it can be appreciated that a stationary peripheral LAN base station may not always be an option when all of the peripheral LAN participants are mobile. In such cases, a high power transmission to reach the
20 premises LAN may be required.

Fig. 28b illustrates other features of the present invention in the use of spontaneous LANs in association with a vehicle which illustrate the capability of automatically establishing a premises and a peripheral
25 LAN when moving in and out of range to perform services and report on services rendered. In particular, like the forklift 3014 of Fig. 28a, a delivery truck 3033 provides

- 113 -

a focal point for a spontaneous LAN utilization. Within the truck 3033, a storage terminal 3031 is docked so as to draw power from the truck 3033's battery supply. Similarly, a computer terminal 3007 may either be docked or ported. Because of greater battery access, the storage terminal 3031 need only be configured for multiple participation in the premises, peripheral and vehicular LANs and in a radio WAN, such as RAM Mobile Data, CDPD, MTEL, ARDIS, etc. The storage terminal 3031, although also capable of premises and peripheral LAN participation, need only be configured for vehicular LAN participation.

Prior to making a delivery, the truck enters a docking area for loading. As goods are loaded into the truck, the information regarding the goods is down-loaded into the storage terminal 3031 via the terminal 3007 or code reader 3009 (Fig. 28a) via the premises or peripheral LAN communications. This loading might also be accomplished automatically as the forklift 3014 comes into range of the delivery truck 3033, establishes or joins the peripheral LAN, and transmits the previously collected data as described above in relation to Fig 28a. Alternately, loading might also be accomplished via the premises LAN.

As information regarding a good is received and stored, the storage terminal 3031 might also request further information regarding any or all of the goods via

- 114 -

the peripheral LAN's link to the host computer 3011 through the premises LAN. More likely however, the storage terminal 3031 if appropriately configured would participate on the premises LAN to communicate directly with the host computer 3011 to retrieve such information.

The peripheral LAN base station 3021 if located on the dock could provide a direct low power peripheral LAN connection to the backbone LAN 3019 and to the host computer 3011. Once fully loaded and prior to leaving the dock, the storage device 3031 may generate a printout of the information relating to the loaded goods via a peripheral LAN established with the printer 3013 on the dock. In addition, the information may be transmitted via the peripheral LAN modem 3023 to a given destination site.

As illustrated in Fig. 28c, once the storage terminal 3031 and hand-held terminal 3007 moves out of range of the premises and peripheral LANs, i.e., the truck 3033 drives away from the dock, the vehicular LAN can only gain access to the premises LAN via the more costly radio WAN communication. Thus, although the storage terminal 3031 might only be configured with relaying control point functionality, to minimize radio WAN communication, the storage terminal 3031 can be configured to store relatively large amounts of information and to provide processing power. Thus, the terminal 3007 can access such information and processing

- 115 -

power without having to access devices on the premises LAN via the radio WAN.

Upon reaching the destination, the storage terminal 3031 may participate in any in range peripheral and premises LAN at the delivery site dock. Specifically, as
5 specific goods are unloaded, they are scanned for delivery verification, preventing delivery of unwanted goods. The driver is also informed if goods that should have been delivered are still in the truck. As this
10 process takes place, a report can also be generated via a peripheral or premises LAN printer at the destination dock for receipt signature. Similarly, the peripheral LAN modem on the destination dock can relay the delivery information back to the host computer 3011 for billing
15 information or gather additional information needed, avoiding use of the radio WAN.

If the truck 3033 is used for service purposes, the truck 3033 leaves the dock in the morning with the addresses and directions of the service destinations,
20 technical manuals, and service notes which have been selectively downloaded from the host computer 3011 via either the premises or peripheral LAN to the storage terminal 3031 which may be configured with a hard drive and substantial processing power. Upon pulling out of
25 range, the storage terminal 3031 and the computer terminal 3007 automatically form an independent, detached vehicular LAN. Alternately, the terminals 3007 and 3031

✓

- 116 -

may have previously formed the vehicular LAN before leaving dock. In one embodiment, the vehicular LAN operates using frequency hopping protocol much the same as that of the premises LAN, with the storage terminal 3031 acting much like the premises LAN base stations. Thus, the radio transceiver circuitry for the premises LAN participation may also be used for the vehicular LAN and, as detailed above, a peripheral LAN. Similarly, if the radio WAN chosen has similar characteristics, it may to be incorporated into a single radio transceiver.

At each service address, the driver collects information using the terminal 3007 either as the data is collected, if within vehicular LAN transmission range of the storage terminal 3031, or as soon as the terminal 3007 comes within range. Any stored information within storage terminal 3031 may be requested via the vehicular LAN by the hand-held terminal 3007. Information not stored within the vehicular LAN may be communicated via a radio WAN as described above.

Referring again to Fig. 28b, upon returning to the dock, the storage terminal 3031, also referred to herein as a vehicle terminal, joins in or establishes a peripheral LAN with the peripheral LAN devices on the dock, if necessary. Communication is also established via the premises LAN. Thereafter, the storage terminal 3031 automatically transfers the service information to the host computer 3011 which uses the information for

- 117 -

billing and in formulating service destinations for automatic downloading the next day.

Fig. 29 is a diagrammatic illustration of another embodiment using a peripheral LAN to supporting roaming data collection by an operator according to the present invention. As an operator 3061 roams the warehouse floor he carries with him a peripheral LAN comprising the terminal 3007, code reader 3009 and a portable printer 3058. The operator collects information regarding goods, such as the box 3010, with the code reader 3009 and the terminal 3007. If the power resources are equal, the terminal 3007 may be configured and designated to also participate in the premises LAN.

Corresponding information to the code data must be retrieved from the host computer 3011. The collected code information and retrieved corresponding information can be displayed on the terminal 3007. After viewing for verification, the information can be printed on the printer 3058. Because of this data flow requirement, the computer terminal 3007 is selected as the peripheral LAN device which must also carry the responsibility of communicating with the premises LAN.

If during collection, the operator decides to power down the computer terminal 3007 because it is not needed, the peripheral LAN becomes detached from the premises LAN. Although it might be possible for the detached peripheral LAN to function, all communication with the

- 118 -

host computer 3011 through the premises LAN is placed in a queue awaiting reattachment. As soon as the detached peripheral LAN comes within range of an attached peripheral LAN device, i.e., a device attached to the premises LAN, the queued communications are relayed to the host. It should be clear from this description that the peripheral LAN may roam in relation to a device attached to the premises LAN ("premises LAN device"). Similarly, the premises LAN device may roam in relation to the peripheral LAN. The roaming constitutes a relative positioning. Moreover, whenever a peripheral LAN and a master device move out of range of each other, the peripheral LAN may either poll for or scan for another master device for attachment. The master device may constitute a premises LAN device, yet need not be.

To avoid detachment when the terminal 3007 is powered down, the code reader 3009 may be designated as a backup to the terminal 3007 for performing the higher power communication with the premises LAN. As described in more detail below in reference to Fig. 33c regarding the idle sense protocol, whenever the code reader 3009 determines that the terminal 3007 has stopped providing access to the premises LAN, the code reader 3009 will take over the role if it is next in line to perform the backup service. Thereafter, when the computer terminal 3007 is powered up, it monitors the peripheral LAN channel, requests and regains from the code reader 3009

- 119 -

the role of providing an interface with the premises LAN. This, however, does not restrict the code reader 3009 from accessing the premises LAN although the reader 3009 may choose to use the computer terminal 3007 for power conservation reasons.

In addition, if the computer terminal 3007 reaches a predetermined low battery threshold level, the terminal 3007 will attempt to pass the burden of providing premises LAN access to other peripheral LAN backup devices. If no backup device exists in the current peripheral LAN, the computer terminal 3007 may refuse all high power transmissions to the premises LAN. Alternatively, the computer terminal 3007 may either refuse predetermined select types of requests, or prompt the operator before performing any transmission to the premises LAN. However, the computer terminal 3007 may still listen to the communications from the premises LAN and inform peripheral LAN members of waiting messages.

Fig. 30 is a block diagram illustrating the functionality of RF transceivers built in accordance with the present invention. Although preferably plugging into PCMCIA slots of the computer terminals and peripherals, the transceiver 3110 may also be built-in or externally attached via available serial, parallel or ethernet connectors for example. Although the transceivers used by potential peripheral LAN master devices may vary from those used by peripheral LAN slave devices (as detailed

- 120 -

below), they all contain the illustrated functional blocks.

In particular, the transceiver 3110 contains a radio unit 3112 which attaches to an attached antenna 3113.

5 The radio unit 3112 used in peripheral LAN slave devices need only provide reliable low power transmissions, and are designed to conserve cost, weight and size. Potential peripheral LAN master devices not only require the ability to communicate with peripheral LAN slave
10 devices, but also require higher power radios to also communicate with the premises LAN. Thus, potential peripheral LAN master devices and other non-peripheral LAN slave devices might contain two radio units 3112 or two transceivers 3110 -- one serving the premises LAN and
15 the other serving the peripheral LAN -- else only contain a single radio unit to service both networks.

In embodiments where cost and additional weight is not an issue, a dual radio unit configuration for potential peripheral LAN master devices may provide
20 several advantages. For example, simultaneous transceiver operation is possible by choosing a different operating band for each radio. In such embodiments, a 2.4GHz radio is included for premises LAN communication while a 27 MHz radio supports the peripheral LAN.

25 Peripheral LAN slave devices receive only the 27 MHz radio, while the non-potential peripheral LAN participants from the premises LAN are fitted with only

- 121 -

the 2.4GHz radios. Potential peripheral LAN master devices receive both radios. The low power 27 MHz peripheral LAN radio is capable of reliably transferring information at a range of approximately 40 to 100 feet asynchronously at 19.2 Kbps. An additional benefit of using the 27 MHz frequency is that it is an unlicensed frequency band. The 2.4 GHz radio provides sufficient power (up to 1 Watt) to communicate with other premises LAN devices. Another benefit of choosing 2.4 GHz or 27 MHz bands is that neither require FCC licensing. Many different frequency choices could also be made such as the 900 MHz band, etc.

In embodiments where cost and additional weight are at issue, a single radio unit configuration is used for potential peripheral LAN master devices. Specifically, in such embodiments, a dual mode 2.4 GHz radio supports both the peripheral LAN and premises LANs. In a peripheral LAN mode, the 2.4 GHz radio operates at a single frequency, low power level (sub-milliwatt) to support peripheral LAN communication at relatively close distances (20-30 feet). In a high power (up to 1 Watt) or main mode, the 2.4 GHz radio provides for frequency-hopping communication over relatively long distance communication connectivity with the premises LAN. Although all network devices might be fitted with such a dual mode radio, only peripheral LAN master devices use both modes. Peripheral LAN slave devices would only use

- 122 -

the low power mode while all other premises LAN devices would use only the high power mode. Because of this, to save cost, peripheral LAN slave devices are fitted with a single mode radio operating in the peripheral LAN mode.

5 Non-peripheral LAN participants are also fitted with a single mode (main mode) radio unit for cost savings.

Connected between the radio unit 3112 and an interface 3110, a microprocessor 3120 controls the information flow between through the transceiver 3110.

10 Specifically, the interface 3115 connects the transceiver 3110 to a selected computer terminal, a peripheral device or other network device. Many different interfaces 3115 are used and the choice will depend upon the connection port of the device to which the transceiver 3110 will be

15 attached. Virtually any type of interface 3110 could be adapted for use with the transceiver 3110 of the present invention. Common industry interface standards include RS-232, RS-422, RS-485, 10BASE2 Ethernet, 10BASE5 Ethernet, 10BASE-T Ethernet, fiber optics, IBM 4/16 Token

20 Ring, V.11, V.24, V.35, Apple Localtalk and telephone interfaces. In addition, via the interface 3115, the microprocessor 3120 maintains a radio independent, interface protocol with the attached network device, isolating the attached device from the variations in

25 radios being used.

The microprocessor 3120 also controls the radio unit 3112 to accommodate communication with the either the

- 123 -

premises LAN, the peripheral LAN, or both (for dual mode radios). Moreover, the same radio might also be used for vehicular LAN and radio WAN communication as described above. More specifically, in a main mode transceiver, the microprocessor 3120 utilizes a premises LAN protocol to communicate with the premises LAN. Similarly, in a peripheral LAN mode transceiver, the microprocessor 3120 operates pursuant to a peripheral LAN protocol to communicate in the peripheral LAN. In the dual mode transceiver, the microprocessor 3120 manages the use of and potential conflicts between both the premises and peripheral LAN protocols. Detail regarding the premises and peripheral LAN protocols can be found in reference to Figs. 33-36 below.

In addition, as directed by the corresponding communication protocol, the microprocessor 3120 controls the power consumption of the radio 3112, itself and the interface 3115 for power conservation. This is accomplished in two ways. First, the peripheral LAN and premises protocols are designed to provide for a low power mode or sleep mode during periods when no communication involving the subject transmitter is desired as described below in relation to Figs. 33-34. Second, both protocols are designed to adapt in both data rate and transmission power based on power supply (i.e., battery) parameters and range information as described in reference to Figs. 35-36.

- 124 -

In order to insure that the proper device is receiving the information transmitted, each device is assigned a unique address. Specifically, the transceiver 3110 can either have a unique address of its own or can use the unique address of the device to which it is attached. The unique address of the transceiver can either be one selected by the operator or system designer or one which is permanently assigned at the factory such as an IEEE address. The address 3121 of the particular transceiver 3110 is stored with the microprocessor 3120.

In the illustrated embodiments of Figs. 28-29, the peripheral LAN master device is shown as being either a peripheral LAN base station or a mobile or portable computer terminal. From a data flow viewpoint, in considering the fastest access through the network, such choices for the peripheral LAN master devices appear optimal. However, any peripheral LAN device might be assigned the role of the master, even those that do not seem to provide an optimal data flow pathway but may provide for optimal battery usage. For example, in the personal peripheral LAN of Fig. 29, because of the support from the belt 3059, the printer might contain the greatest battery capacity of the personal peripheral LAN devices. As such, the printer might be designated the peripheral LAN master device and be fitted with either a dual mode radio or two radios as master devices require. The printer, or other peripheral LAN slave devices, might

- 125 -

also be fitted with such required radios to serve only as a peripheral LAN master backup. If the battery power on the actual peripheral LAN master, i.e., the hand-held terminal 3007 (Fig. 29, drops below a preset threshold, the backup master takes over.

Fig. 31 is a drawing which illustrates an embodiment of the personal peripheral LAN shown in Fig. 29 which designates a printer as the peripheral LAN master device. Specifically, in a personal peripheral LAN 3165, a computer terminal 3170 is strapped to the forearm of the operator. A code reader 3171 straps to the back of the hand of the user and is triggered by pressing a button 3173 with the thumb. Because of their relatively low battery energy, the computer terminal 3170 and code reader 3171 are designated peripheral LAN slave devices and each contain a peripheral LAN transceiver having a broadcast range of two meters or less. Because of its greater battery energy, the printer 3172 contains a dual mode radio and is designated the peripheral LAN master device.

Fig. 32 is a block diagram illustrating a channel access algorithm used by peripheral LAN slave devices according to the present invention. At a block 3181, when a slave device has a message to send, it waits for an idle sense message to be received from the peripheral LAN master device at a block 3183. When an idle sense message is received, the slave device executes a back-off

- 126 -

protocol at a block 3187 in an attempt to avoid collisions with other slave devices waiting to transmit. Basically, instead of permitting every slave device from repeatedly transmitting immediately after an idle sense message is received, each waiting slave is required to first wait for a pseudo-random time period before attempting a transmission. The pseudo-random back-off time period is generated and the waiting takes place at a block 3187. At a block 3189, the channel is sensed to determine whether it is clear for transmission. If not, a branch is made back to the block 3183 to attempt a transmission upon receipt of the next idle sense message. If the channel is still clear, at a block 3191, a relatively small "request to send" type packet is transmitted indicating the desire to send a message. If no responsive "clear to send" type message is received from the master device, the slave device assumes that a collision occurred at a block 3193 and branches back to the block 3183 to try again. If the "clear to send" message is received, the slave device transmits the message at a block 3195.

Several alternate channel access strategies have been developed for carrier sense multiple access (CSMA) systems and include 1-persistent, non-persistent and p-persistent. Such strategies or variations thereof could easily be adapted to work with the present invention.

- 127 -

Fig. 33a is a timing diagram of the protocol used according to one embodiment the present invention illustrating a typical communication exchange between a peripheral LAN master device having virtually unlimited power resources and a peripheral LAN slave device. Time line 3201 represents communication activity by the peripheral LAN master device while time line 3203 represents the corresponding activity by the peripheral LAN slave device. The master periodically transmits an idle sense message 3205 indicating that it is available for communication or that it has data for transmission to a slave device. Because the master has virtually unlimited power resources, it "stays awake" for the entire time period 3207 between the idle sense messages 3205. In other words, the master does not enter a power conserving mode during the time periods 3207.

The slave device uses a binding protocol (discussed below with regard to Fig. 33c) to synchronize to the master device so that the slave may enter a power conserving mode and still monitor the idle sense messages of the master to determine if the master requires servicing. For example, referring to Fig. 33a, the slave device monitors an idle sense message of the master during a time period 3209, determines that no servicing is required, and enters a power conserving mode during the time period 3211. The slave then activates during a time period 3213 to monitor the next idle sense message

- 128 -

of the master. Again, the slave determines that no servicing is required and enters a power conserving mode during a time period 3215. When the slave activates again during a time period 3217 to monitor the next idle sense message, it determines from a "request to send" type message from the master that the master has data for transmission to the slave. The slave responds by sending a "clear to send" type message during the time period 3217 and stays activated in order to receive transmission of the data. The master is thus able to transmit the data to the slave during a time period 3219. Once the data is received by the slave at the end of the time period 3221, the slave again enters a power conserving mode during a time period 3223 and activates again during the time period 3225 to monitor the next idle sense message.

Alternatively, the slave may have data for transfer to the master. If so, the slave indicates as such to the master by transmitting a message during the time period 3217 and then executes a backoff algorithm to determine how long it must wait before transmitting the data. The slave determines from the backoff algorithm that it must wait the time period 3227 before transmitting the data during the time period 3221. The slave devices use the backoff algorithm in an attempt to avoid the collision of data with that from other slave devices which are also trying to communicate with the master. The backoff

- 129 -

algorithm is discussed more fully above in reference to Fig. 32.

The idle sense messages of the master may also aid in scheduling communication between two slave devices. For example, if a first slave device has data for transfer to a second slave device, the first slave sends a message to the master during the time period 3209 requesting communication with the second slave. The master then broadcasts the request during the next idle sense message. Because the second slave is monitoring the idle sense message, the second slave receives the request and stays activated at the end of the idle sense message in order to receive the communication. Likewise, because the first slave is also monitoring the idle sense message, it too receives the request and stays activated during the time period 3215 to send the communication.

Fig. 33b is a timing diagram of the protocol used according to one embodiment of the present invention illustrating a typical communication exchange between a peripheral LAN master having limited power resources and a peripheral LAN slave device. This exchange is similar to that illustrated in Fig. 33a except that, because it has limited power resources, the master enters a power conserving mode. Before transmitting an idle sense message, the master listens to determine if the channel is idle. If the channel is idle, the master transmits an idle sense message 3205 and then waits a time period 3231

- 130 -

to determine if any devices desire communication. If no communication is desired, the master enters a power conserving mode during a time period 3233 before activating again to listen to the channel. If the channel is not idle, the master does not send the idle sense message and enters a power saving mode for a time period 3235 before activating again to listen to the channel.

Communication between the master and slave devices is the same as that discussed above in reference to Fig 33a except that, after sending or receiving data during the time period 3219, the master device enters a power conserving mode during the time period 3237.

Fig. 33c is also a timing diagram of one embodiment of the protocol used according to the present invention which illustrates a scenario wherein the peripheral LAN master device fails to service peripheral LAN slave devices. The master device periodically sends an idle sense message 3205, waits a time period 3231, and enters a power conserving mode during a time period 3233 as discussed above in reference to Fig 33b. Similarly, the slave device monitors the idle sense messages during time periods 3209 and 3213 and enters a power conserving mode during time periods 3211 and 3215. For some reason, however, the master stops transmitting idle sense messages. Such a situation may occur, for example, if the master device is portable and is carried outside the range of the slave's radio. During a time period 3241,

- 131 -

the slave unsuccessfully attempts to monitor an idle sense message. The slave then goes to sleep for a time period 3243 and activates to attempt to monitor a next idle sense message during a time period 3245, but is again unsuccessful.

The slave device thereafter initiates a binding protocol to attempt to regain synchronization with the master. While two time periods 3241 and 3245 are shown, the slave may initiate such a protocol after any number of unsuccessful attempts to locate an idle sense message. With this protocol, the slave stays active for a time period 3247, which is equal to the time period from one idle sense message to the next, in an attempt to locate a next idle sense message. If the slave is again unsuccessful, it may stay active until it locates an idle sense message from the master, or, if power consumption is a concern, the slave may enter a power conserving mode at the end of the time period 3247 and activate at a later time to monitor for an idle sense message.

In the event the master device remains outside the range of the slave devices in the peripheral LAN for a period long enough such that communication is hindered, one of the slave devices may take over the functionality of the master device. Such a situation is useful when the slave devices need to communicate with each other in the absence of the master. Preferably, such a backup device has the ability to communicate with devices on the

- 132 -

premises LAN. If the original master returns, it listens to the channel to determine idle sense messages from the backup, indicates to the backup that it has returned and then begins idle sense transmissions when it
5 reestablishes dominance over the peripheral LAN.

Fig. 34 is a timing diagram illustrating one embodiment of the peripheral LAN master device's servicing of both the high powered premises LAN and the low powered peripheral LAN subnetwork, with a single or
10 plural radio transceivers, in accordance with present invention. Block 3251 represents typical communication activity of the master device. Line 3253 illustrates the master's communication with a base station on the premises LAN while line 3255 illustrates the master's
15 communication with a slave device on the peripheral LAN. Lines 3257 and 3259 illustrate corresponding communication by the base station and slave device, respectively.

The base station periodically broadcasts HELLO
20 messages 3261 indicating that it is available for communication. The master device monitors the HELLO messages during a time period 3263, and, upon determining that the base does not need servicing, enters a power conserving mode during a time period 3265. The
25 master then activates for a time period to monitor the next HELLO message from the base. If the master has data to send to the base, it transmits the data during a time

- 133 -

period 3271. Likewise, if the base has data to send to the master, the base transmits the data during a time period 3269. Once the data is received or sent by the master, it may again enter a power conserving mode.

5 While HELLO message protocol is discussed, a number of communication protocols may be used for communication between the base and the master device. As may be appreciated, the peripheral LAN master device acts as a slave to base stations in the premises LAN.

10 Generally, the communication exchange between the master and the slave is similar to that described above in reference to Fig. 33b. Block 3273, however, illustrates a situation where the master encounters a communication conflict, i.e., it has data to send to or
15 receive from the slave on the peripheral LAN at the same time it will monitor the premises LAN for HELLO messages from the base. If the master has two radio transceivers, the master can service both networks. If, however, the master only has one radio transceiver, the master chooses
20 to service one network based on network priority considerations. For example, in block 3273, it may be desirable to service the slave because of the presence of data rather than monitor the premises LAN for HELLO messages from the base. On the other hand, in block
25 3275, it may be more desirable to monitor the premises LAN for HELLO messages rather than transmit an idle sense message on the peripheral LAN.

- 134 -

Figs. 35 and 36 are block diagrams illustrating additional power saving features according to the present invention, wherein ranging and battery parameters are used to optimally select the appropriate data rate and power level for subsequent transmissions. Specifically, even though network devices such as the computer terminal 3007 in Figs. 28-29 have the capability of performing high power transmissions, because of battery power concerns, such devices are configured to utilize minimum transmission energy. Adjustments are made based on ranging information and on battery parameters. Similarly, within the peripheral LAN, even though lower power transceivers are used, battery conservation issues also justify the use of such data rate and power adjustments. This process is described in more detail below in reference to Figs. 35 and 36.

More specifically, Fig. 35 is a block diagram which illustrates a protocol 3301 used by a destination peripheral LAN device and a corresponding protocol 3303 used by a source peripheral LAN device to adjust the data rate and possibly the power level for future transmission between the two devices. At a block 3311, upon receiving a transmission from a source device, the destination device identifies a range value at a block 3313. In a low cost embodiment, the range value is identified by considering the received signal strength indications (RSSI) of the incoming transmission. Although RSSI

- 135 -

circuitry might be placed in all peripheral LAN radios, the added expense may require that only peripheral LAN master devices receive the circuitry. This would mean that only peripheral LAN master devices would perform the function of the destination device. Other ranging techniques or signal quality assessments can also be used, such as measuring jitter in received signals, by adding additional functionality to the radios. Finally, after identifying the range value at the block 3313, the destination device subsequently transmits the range value to the slave device from which the transmission was received, at a block 3314.

Upon receipt of the range value from the destination device at a block 3321, the source peripheral LAN device evaluates its battery parameters to identify a subsequent data rate for transmission at a block 3323. If range value indicates that the destination peripheral LAN device is very near, the source peripheral LAN device selects a faster data rate. When the range value indicates a distant master, the source device selects a slower rate. In this way, even without adjusting the power level, the total energy dissipated can be controlled to utilize only that necessary to carry out the transmission. However, if constraints are placed on the maximum or minimum data rates, the transmission power may also need to be modified. For example, to further minimize the complexity associated with a fully random

- 136 -

range of data rate values, a standard range and set of several data rates may be used. Under such a scenario, a transmission power adjustment might also need to supplement the data rate adjustment. Similarly, any
5 adjustment of power must take into consideration maximum and minimum operable levels. Data rate adjustment may supplement such limitations. Any attempted modification of the power and data rate might take into consideration any available battery parameters such as those that might
10 indicate a normal or current battery capacity, the drain on the battery under normal conditions and during transmission, or the fact that the battery is currently being charged. The latter parameter proves to be very significant in that when the battery is being charged,
15 the peripheral LAN slave device has access to a much greater power source for transmission, which may justify the highest power transmission and possibly the slowest data rate under certain circumstances.

Finally, at a block 3325, an indication of the
20 identified data rate is transmitted back to the destination device so that future transmissions may take place at the newly selected rate. The indication of data rate may be explicit in that a message is transmitted designating the specific rate. Alternately, the data
25 rate may be transferred implicitly in that the new rate is chose and used by the source, requiring the

- 137 -

destination to adapt to the change. This might also be done using a predefined header for synchronization.

In addition, at the block 3325, in another embodiment, along with the indication of the identified data rate, priority indications are also be communicated. Whenever battery power is detected as being low, a radio transmits a higher priority indication, and each receiver thereafter treats the radio as having a higher protocol priority than other such radios that exhibit normal power supply energy. Thus, the remaining battery life is optimized. For example, in a non-polling network, the low power device might be directly polled periodically so to allow scheduled wake-ups and contention free access to a receiver. Similarly, in an alternate embodiment, priority indications not need to be sent. Instead, the low battery power device itself exercises protocol priority. For example, for channel access after detecting that the channel is clear at the end of an ongoing transmission, devices with normal energy levels are required to undergo a pseudo-random back-off before attempting a transmission (to avoid collision). The low power device may either minimize the back-off period or ignore the back-off period completely. Thus, the low power device gains channel access easier than other normal power level devices. Other protocol priority schemes may also be assigned by the receivers to the low

- 138 -

power device (via the indication), else may be taken directly by the low power device.

Fig. 36 illustrates an alternate embodiment for carrying out the data rate and possibly power level adjustment. At a block 3351 upon binding and possibly periodically, the source peripheral LAN device sends an indication of its current battery parameters to the destination peripheral LAN device. This indication may be each of the parameters or may be an averaged indication of all of the parameters together. At a block 3355, upon receipt, the destination peripheral LAN device stores the battery parameters (or indication). Finally, at a block 3358, upon receiving a transmission from the source device, based on range determinations and the stored battery parameters, the destination terminal identifies the subsequent data rate (and possibly power level). Thereafter, the new data rate and power level are communicated to the source device either explicitly or implicitly for future transmissions.

Fig. 37 illustrates an exemplary block diagram of a radio unit 3501 capable of concurrent participation on multiple LAN's. To transmit, a control processor 3503 sends a digital data stream to a modulation encoding circuit 3505. The modulation encoding circuit 3505 encodes the data stream in preparation for modulation by frequency translation circuit 3507. The carrier frequency used to translate the data stream is provided

- 139 -

by a frequency generator circuit 3509. Thereafter, the modulated data stream is amplified by a transmitter amplifier circuit 3511 and then radiated via the one of a plurality of antennas 3513 that has been selected via
5 an antenna switching circuit 3515. Together, the modulation encoding circuitry 3505, translator 3507, amplifier 3511 and associated support circuitry constitute the transmitter circuitry.

Similarly, to receive data, the RF signal received
10 by the selected one of the plurality of antennas 3513 is communicated to a receiver RF processing circuit 3517. After performing a rather coarse frequency selection, the receiver RF processing circuit 3517 amplifies the RF signal received. The amplified received signal undergoes
15 a frequency shift to an IF range via a frequency translation circuit 3519. The frequency translation circuit 3519 provides the center frequency for the frequency shift. Thereafter, a receiver signal processing circuit receives the IF signal, performs a
20 more exact channel filtering and demodulation, and forwards the received data to the control processor 3503, ending the process. Together, the receiver signal processing 3521, translator 3517, receiver RF processing 3517 and associated support circuitry constitute the
25 receiver circuitry.

The control processor 3503 operates pursuant to a set of software routines stored in memory 3522 which may

- 140 -

also store incoming and outgoing data. Specifically, the memory 3522 contains routines which define a series of protocols for concurrent communication on a plurality of LANs. As part of such operation, the control processor 3503 provides for power savings via a power source control circuit 3523, i.e., whenever the participating protocols permit, the control processor 3503 causes selective power down of the radio transceiver circuitry via a control bus 3525. Also via the bus 3525, the control processor sets the frequency of the frequency generator 3509 so as to select the appropriate band and channel of operation required by a correspondingly selected protocol. Similarly, the control processor 3503 selects the appropriate antenna (via the antenna switching circuitry 3515) and channel filtering in preparation for operation on a selected LAN.

Responding to the software routines stored in the memory 3522, the control processor 3503 selects the appropriate LANs to establish participation, detaches from those of the selected LANs in which participation is no longer needed, identifies from the selected LANs a current priority LAN in which to actively participate, maintains a time-shared servicing of the participating LANs. Further detail regarding this process follows below.

In one embodiment, the control processor 3503 constitutes a typical microprocessor on an independent integrated circuit. In another embodiment, the control

- 141 -

processor 3503 comprises a combination of distributed processing circuitry which could be included in a single integrated circuit as is a typical microprocessor. Similarly, the memory 3522 could be any type of memory unit(s) or device(s) capable of software storage.

The radio circuitry illustrated is designed with the frequency nimble frequency generator 3509 so as to be capable of operation on a plurality of LANs/WANs. Because each of the plurality may be allocated different frequency bands, more than one antenna may be desirable (although a single antenna could be used, antenna bandwidth limitations might result in an unacceptable transmission-reception inefficiency). Thus, to select the appropriate configuration, the control processor 3503 first identifies the LAN/WAN on which to participate and selects the corresponding radio configuration parameters from the memory 3521. Thereafter, using the configuration parameters and pursuant to control routines stored in the memory 3522, the control processor 3503 sets the frequency of the generator 3509, selects the appropriate antenna via the antenna switching circuit 3515, and configures the receiver RF and signal processing circuits 3517 and 3521 for the desired LAN/WAN.

More particularly, the antenna switching circuit 3515 comprises a plurality of digitally controlled switches, each of which is associated with one of the

- 142 -

plurality of antennas 3513 so as to permit selective connection by the control processor 3503 of any available antenna to the transceiver circuitry.

Fig. 38 illustrates an exemplary functional layout of the frequency generator 3509 of Fig. 37 according to one embodiment of the present invention. Basically, the frequency generator 3509 responds to the control processor 3503 by producing the translation frequency necessary for a selected LAN/WAN. The illustrated frequency generator comprises a voltage controlled oscillator (VCO) 3601. As is commonly known, for a VCO, the center frequency F_{VCO} tracks the input voltage. However, because typical VCO's are subject to drift, the VCO is stabilized by connecting it in a phase locked loop to a narrowband reference, such as a crystal reference oscillator 3603. The oscillator 3603 outputs a signal of a fixed or reference frequency F_{REF} to a divide-by-R circuit 3605, which divides as its name implies the reference frequency F_{REF} by the known number R. A phase detector 3609 receives the divided-by-R output of the circuit 3605 and the feedback from the output of the VCO 3601 via a divide-by-N circuit 3607. Upon receipt, the phase detector 3609 compares the phase of the outputs from the circuits 3605 and 3607. Based on the comparison, a phase error signal is generated and applied to a low-pass loop filter 3611. The output of the filter 3611 is applied to the input of the VCO 3601 causing the

- 143 -

center frequency of the VCO 3601 to lock-in. Therefore,
if the output of the VCO 3601 begins to drift out of
phase of the reference frequency, the phase detector 3609
responds with a corrective output so as to adjust the
5 center frequency of the VCO 3601 back in phase.

With the illustrated configuration, the center
frequency of the VCO 3601 is a function of the reference
frequency as follows:

$$F_{VCO} = (F_{REF} * N) / R$$

10 Thus, to vary the center frequency of the VCO 3601 to
correspond to a band of a selected LAN/WAN in which
active participation is desired, the control processor
3503 (Fig. 37) need only vary the variables "R" and "N"
and perhaps the frequency of the reference oscillator.
15 Because the output F_{REF} of the reference oscillator 3603
is quite stable, the phase lock loop as shown also keeps
the output frequency F_{VCO} of the VCO 3601 stable.

More specifically, although any other scheme might
be implemented, the value R in the divide-by-R circuit
20 3605 is chosen so as to generate an output equal to the
channel spacing of a desired LAN/WAN, while the value N
is selected as a multiplying factor for stepping up the
center frequency of the VCO 3601 to the actual frequency
of a given channel. Moreover, the frequency of the
25 reference oscillator is chosen so as to be divisible by
values of R to yield the channel spacing frequencies of
all potential LANs and WANs. For example, to participate

- 144 -

on both MTEL Corporation's Two Way Paging WAN (operating at 900MHz with 25KHz and 50KHz channel spacings) and ARDIS Corporation's 800MHz specialized mobile radio (SMR) WAN (operating at 25KHz channel spacings centered at multiples of 12.5KHz), a single reference frequency may be chosen to be a whole multiple of 12.5KHz. Alternately, multiple reference frequencies may be chosen. Moreover, the value N is chosen to effectively multiply the output of the divide-by-R circuit 3605 to the base frequency of a given channel in the selected WAN.

For frequency hopping protocols, the value of R is chosen so as to yield the spacing between frequency hops. Thus, as N is incremented, each hopping frequency can be selected. Randomizing the sequence of such values of N provides a hopping sequence for use by a base station as described above. Pluralities of hopping sequences (values of N) may be stored in the memory 3522 (Fig. 37) for operation on the premises LAN, for example.

In addition to the single port phase locked loop configuration for the frequency generator 3509, other configurations might also be implemented. Exemplary circuitry for such configurations can be found in copending U.S. Patent Application Ser. No. _____ (Attorney Docket Nos. DN37139XXA; 10458US03), filed March 4, 1994 by Mahany et al., entitled "Method of and Apparatus For Controlling Modulation of Digital Signals

- 145 -

in Frequency-Modulated Transmissions". This application is incorporated herein in its entirety.

Fig. 39 illustrates further detail of the receiver RF processing circuit 3517 of Fig. 37 according to one embodiment of the present invention. Specifically, a preselector 3651 receives an incoming RF data signal from a selected one of the plurality of antennas 3513 (Fig. 37) via an input line 3653. The preselector 3651 provides a bank of passive filters 3657, such as ceramic or dielectric resonator filters, each of which provides a coarse filtering for one of the LAN/WAN frequencies to which it is tuned. One of the outputs from the bank of passive filters 3657 is selected by the control processor 3503 via a switching circuit 3655 so as to monitor the desired one of the available LANs/WANs. Thereafter, the selected LAN/WAN RF signal is amplified by an RF amplifier 3659 before translation by the frequency translation circuit 3519 (Fig. 37).

Fig. 40 illustrates further detail of the receiver signal processing circuit 3521 of Fig. 37 according to one embodiment of the present invention. In particular, digitally controlled switching circuits 3701 and 3703 respond to the control processor 3503 by selecting an appropriate pathway for the translated IF data signal through one of a bank of IF filters 3705. Each IF filter is an analog crystal filter, although other types of filters such as a saw filter might be used. The IF

- 146 -

filters 3705 provide rather precise tuning to select the specific channel of a given LAN/WAN.

After passing through the switching circuit 3703, the filtered IF data signal is then amplified by an IF amplifier 3707. The amplified IF signal is then communicated to a demodulator 3709 for demodulation. The control processor retrieves the incoming demodulated data signal for processing and potential storage in the memory 3522 (Fig. 37).

Fig. 41 illustrates further detail of the receiver signal processing circuit 3521 of Fig. 37 according to another embodiment of the present invention. Specifically, the IF signal resulting from the translation by the frequency translator circuitry 3519, enters the receiver signal processing circuit via an input 3751. Thereafter, the IF signal passes through an anti-aliasing filter 3753, and is amplified by a linear amplifier 3755. An IF oscillator 3757 supplies a reference signal f_{REF} for translation of the incoming IF signal at frequency translation circuits 3759 and 3761. A phase shift circuit 3763 provides for a 90 degree shift of f_{REF} , i.e., if f_{REF} is considered a SINE wave, then the output of the circuit 3763 is the COSINE of f_{REF} . Both the SINE and COSINE frequency translation pathways provide for channel selection of the incoming data signal. Thereafter the data signals are passed through corresponding low pass filters 3765 and 3767 in

- 147 -

preparation for sampling by analog to digital (A/D) converters 3769 and 3771. Each A/D converter forwards the sampled data to a digital signal processor 3773 which provides for further filtering and demodulation. The digital signal processor 3773 thereafter forwards the incoming data signal to the control processor 3503 (Fig. 37) via an output line 3775. Moreover, although the digital signal processor 3773 and the control processor 3507 are discrete components in the illustrated example, they may also be combined into a single integrated circuit.

Fig. 42 illustrates further detail of some of the storage requirements of the memory 3522 of Fig. 37 according to one embodiment of the present invention. To control the radio, the control processor 3503 (Fig. 37) accesses the information in the memory 3522 needed for radio setup and operation on a plurality of LANs/WANs. Among other information, the memory 3522 stores: 1) a plurality of software protocols, one for each LAN/WAN to be supported, which define how the radio is to participate on the corresponding LAN; and 2) an overriding control set of routines which govern the selection, use and interaction of the plurality of protocols for participation on desired LANs/WANs.

Specifically, in the memory unit 3522, among other information and routines, software routines relating to the media access control (MAC) sublayer of the

- 148 -

communication protocol layers can be found. In general, a MAC sublayer provides detail regarding how communication generally flows through a corresponding LAN or WAN. Specifically, the MAC sublayer handles functions such as
5 media access control, acknowledge, error detection and retransmission. The MAC layer is fairly independent of the specific radio circuitry and channel characteristics of the LAN or WAN.

As illustrated, premises LAN, peripheral LAN,
10 vehicular LAN and WAN MAC routines 3811, 3813, 3815 and 3817 provide definition as to how the control processor 3503 (Fig. 37) should operate while actively participating on each LAN or WAN. Although only the several sets of MAC routines are shown, many other sets
15 might also be stored or down-loaded into the memory 3522. Moreover, the sets of MAC routines 3811-17 might also share a set of common routines 3819. In fact, the sets of MAC routines 3811-17 might be considered a subset of an overall MAC which shares the common MAC routines 3819.

20 Below the MAC layer in the communication hierarchy, hardware and channel related software routines and parameters are necessary for radio control. For example, such routines govern the specific switching for channel filtering and antenna selection required by a given LAN
25 or WAN. Similarly, these routines govern the control processor 3503's selection of parameters such as for R and N for the frequency generator 3509 (Fig. 38), or the

- 149 -

selective power-down (via the power source control circuitry 3503 -- Fig. 37) of portions or all of the radio circuitry whenever possible to conserve battery power. As illustrated, such routines and parameters are referred to as physical (PHY) layer control software 3821. Each of the sets of MAC routines 3811-17 and 3819 provide specific interaction with the PHY layer control software 3821.

A set of MAC select/service routines 3823 govern the management of the overall operation of the radio in the network. For example, if participation on the premises LAN is desired, the MAC select/service routines 3823 direct the control processor 3503 (Fig. 37) to the common and premises MAC routines 3819 and 3811 respectively.

Thereafter, if concurrent participation with a peripheral LAN is desired, the select/service routines 3823 direct the control processor 3503 to enter a sleep mode (if available). The control processor 3503 refers to the premises LAN MAC routines 3811, and follows the protocol necessary to establish sleep mode on the premises LAN. Thereafter, the select/service routines 3823 directs the control processor 3503 to the peripheral LAN MAC routines 3813 to establish and begin servicing the peripheral LAN. Whenever the peripheral LAN is no longer needed, the select/service routines 3823 direct a detachment from the peripheral LAN (if required) as specified in the peripheral LAN MAC routines 3813. Similarly, if during

- 150 -

the servicing of the peripheral LAN a overriding need to service the premises LAN arises, the processor 3503 is directed to enter a sleep mode via the peripheral LAN MAC routines 3813, and to return to servicing the premises LAN.

Although not shown, additional protocol layers as well as incoming and outgoing data are also stored with the memory 3522, which, as previously articulated, may be a distributed plurality of storage devices.

Fig. 43 illustrates a software flow chart describing the operation of the control processor 3503 (Fig. 37) in controlling the radio unit to participate on multiple LANs according to one embodiment of the present invention. Specifically, at a block 3901, the control processor first determines whether the radio unit needs to participate on an additional LAN (or WAN). If such additional participation is needed, at a block 3903, the radio unit may register sleep mode operation with other participating LANs if the protocols of those LANs so require and the radio unit has not already done so. Next, at a block 3905, the control processor causes the radio unit to poll or scan to locate the desired additional LAN. If the additional LAN is located at a block 3907, participation of the radio unit on the additional LAN is established at a block 3909.

If additional participation is not needed at block 3901, or if the additional LAN has not been located at

- 151 -

block 3907, or once participation of the radio unit on the additional LAN has been established at block 3909, the control processor next determines at a block 3911 whether any of the participating LANs require servicing.

5 If any given participating LAN requires servicing, , at a block 3913, the radio unit may be required by the protocol of the given LAN to reestablish an active participation status on that LAN, i.e., indicate to the given LAN that the radio unit has ended the sleep mode.

10 Next, at a block 3915, the radio unit services the given LAN as needed or until the servicing of another LAN takes priority over that of the given LAN. At a block 3917, the radio unit may then be required to register sleep mode operation with the given LAN if the LAN's protocol

15 so requires.

At that point, or if no participating LAN needs servicing at block 3911, the control processor determines at a block 3919 whether the radio needs to detach from any given participating LAN. If so, the radio unit may

20 implicitly detach at a block 3923 if the protocol of the LAN from which the radio wishes to detach requires no action by the radio unit. However, at a block 3921, the radio unit may be required to establish active participation on the LAN in order to explicitly detach at

25 block 3923. For example, such a situation may arise when a portable terminal desires to operate on a shorter range vehicular LAN and detaches from a premises LAN. The

- 152 -

portable terminal may be required by the protocol of the premises LAN to establish active communication on the premises LAN to permit the radio unit to inform the premises LAN that it is detaching and can only be
5 accessed through the vehicular LAN.

Once the radio unit is detached at block 3923, or if the radio unit does not need to detach from any participating LANs at block 3919, the control processor returns to block 3901 to again determine whether the
10 radio unit needs to participate on an additional LAN, and repeats the process.

Fig. 44 is an alternate embodiment of the software flow chart wherein the control processor participates on a master LAN and, when needed, on a slave LAN. Specifically, at a block 3951, the control processor
15 causes the radio unit to poll or scan in order to locate the master LAN. If the master LAN has not been located at a block 3953, polling or scanning for the master LAN continues. Once the master LAN is located, participation with the master is established at a block 3955. At a
20 block 3957, the radio unit participates with the master LAN until the need for the radio unit to participate on the slave LAN takes precedence. When that condition occurs, the control processor determines at a block 3959 whether participation of the radio unit on the slave
25 network is established. If not, such participation is established at a block 3961. Next, at a block 3963, the

- 153 -

radio unit services the slave LAN as needed or until the servicing of the master LAN takes priority. If the control processor determines at a block 3965 that servicing of the slave LAN has been completed, the radio unit detaches from the slave LAN at a block 3967 and returns to block 3957 to continue participation on the master LAN.

However, if the control processor determines at block 3965 that servicing has not been, or may not be, completed, the radio unit does not detach from the slave LAN. In that case, before returning to block 3957 to service the master LAN, the radio unit may be required by the protocol of the slave LAN to register sleep mode operation with the slave LAN at a block 3969.

Moreover, it will be apparent to one skilled in the art having read the foregoing that various modifications and variations of this communication system according to the present invention are possible and is intended to include all those which are covered by the appended claims.

- 154 -

CLAIMS

WE CLAIM:

1. A radio unit for operation in a communication system having a plurality of RF communication networks comprising:

5 a transceiver capable of participating on a plurality of RF communication networks;

10 a memory device which stores a plurality of communication protocols, each communication protocol governing radio operation on one of the plurality of RF communication networks;

15 a control processor which selects from the memory device ones of the plurality of communication protocols to enable the transceiver to simultaneously participate on corresponding ones of the plurality of RF communication networks; and

the control processor resolving conflicts among the participating ones of the plurality of communication protocols.

20 2. The radio unit of claim 1 wherein the control processor causes the radio unit to enter a low power consumption mode when the radio unit is not communicating on any of the plurality of RF communication networks.

25 3. The radio unit of claim 1 wherein, upon detecting a low power supply condition, the control

- 155 -

processor causes the selection of a higher priority protocol to conserve energy.

4. The radio unit of claim 1 wherein, responsive to power supply conditions, the control processor adjusts
5 the data rate for subsequent transmissions.

5. The radio unit of claim 1 wherein the control processor resolves conflicts by continually prioritizing participation on the participating ones of the plurality of RF communication networks.

10 6. A radio unit for operation in a communication system having a plurality of wireless communication networks comprising:

a transceiver capable of participating on a plurality of wireless communication networks;

15 a memory device which stores a plurality of communication protocols, each communication protocol governing radio operation on one of the plurality of wireless communication networks;

20 a control processor which selects from the memory device ones of the plurality of communication protocols to enable the transceiver to simultaneously participate on corresponding ones of the plurality of wireless communication networks; and

- 156 -

the control processor causing the radio unit to enter a low power consumption mode when the radio unit is not communicating on the plurality of wireless communication networks.

5 7. The radio unit of claim 6 wherein the low power consumption mode allocates a higher participation priority to the radio unit in the participating ones of the plurality of wireless communication networks.

10 8. The radio unit of claim 6 wherein, responsive to power supply conditions, the control processor adjusts the data rate for subsequent transmissions.

9. A radio unit for operation in a communication system having a plurality of wireless communication networks comprising:

15 a transceiver capable of participating on a plurality of wireless communication networks;

 a memory device which stores a plurality of communication protocols, each communication protocol governing radio operation on one of the plurality of wireless communication networks; and

20

 a control processor which selects from the memory device ones of the plurality of communication protocols to enable the transceiver to simultaneously participate

- 157 -

on corresponding ones of the plurality of wireless communication networks.

10. The radio unit of claim 9 wherein each of the communication protocols comprises means for establishing communication of the transceiver on the corresponding ones of the plurality of communication networks.

11. The radio unit of claim 9 further comprising a power source which is monitored by the control processor, and wherein the control processor automatically initiates power saving protocols in response to predefined power source conditions.

12. A communication system comprising:

a first RF communication network operating using a first communication protocol;

15 a second RF communication network operating using a second communication protocol;

a radio unit having a control processor, a memory device, and a transceiver, the memory device storing the first and second communication protocols and the transceiver capable of participating on both the first and second RF communication networks;

20 the radio unit participating on the first RF communication network pursuant to the first communication protocol while simultaneously participating on the second

- 158 -

RF communication network pursuant to the second communication protocol; and

the control processor resolving conflicts between the first and second communication protocols.

5 13. The communication system of claim 11 wherein the radio unit participates as a slave on the first RF communication network and as a master on the second RF communication network.

10 14. The communication system of claim 11 wherein the control unit assigns and occasionally modifies a participation priority to the first RF communication network or the second RF communication network.

15 15. The communication system of claim 11 wherein the first RF network comprises a premises communication network.

16. The communication system of claim 11 wherein the second RF network comprises a spontaneous communication network.

20 17. The communication system of claim 16 wherein the spontaneous communication network comprises a peripheral communication network.

- 159 -

18. The communication system of claim 16 wherein the spontaneous communication network comprises a vehicular communication network.

19. A communication system comprising:

5

an RF premises communication network;

an RF spontaneous communication network;

10

a radio unit having a control processor and a transceiver capable of communicating on both the RF premises communication network and the RF spontaneous communication network; and

the radio unit participating as a slave in the RF premises network while participating as a master in the RF spontaneous communication network pursuant to control by the control processor.

15

20. The communication system of claim 19 wherein the RF spontaneous network comprises a peripheral communication network.

20

21. The communication system of claim 20 wherein the peripheral communication network uses narrow band, single frequency operation.

22. The communication system of claim 19 wherein the spontaneous communication network comprises a vehicular local area network.

SUBSTITUTE SHEET (RULE 26)

- 160 -

23. The communication system of claim 22 wherein the vehicular local area network uses frequency hopping operation.

5 24. The communication system of claim 19 wherein the RF premises local area network uses frequency hopping operation.

25. A communication system comprising:
a first RF local area network operating using a first communication protocol;
10 a second RF local area network operating using a second communication protocol;
a radio unit having a control processor and a transceiver capable of participating on both the first and second RF local area networks;
15 the radio unit participating on the first RF local area network pursuant to the first communication protocol while participating on the second RF local area network pursuant to the second communication protocol; and
the control processor resolving conflicts between
20 the first and second communication protocols and causing the radio unit to enter a state of low power consumption when the radio unit is not communicating on either the first or the second RF local area networks.

- 161 -

26. The communication system of claim 20 wherein the radio unit participates as a slave on the first local area network and as a master on the second local area network.

5 27. A communication system comprising:
 a first RF local area network comprising a first plurality of network devices;
 a second RF local area network comprising a second plurality of network devices;
10 a radio unit having a control processor and configured to participate as a member of both the first and second pluralities of network devices;
 when within range of one of the second plurality of network devices, the radio unit participates in the
15 second radio network pursuant to control by the control processor; and
 when within range of one of the first plurality of network devices, the radio unit participates in the second radio network pursuant to control by the control
20 processor.

28. The communication system of claim 26 wherein the radio unit further comprises a power supply having conditions of operation, and wherein the control processor responds to the power supply conditions by
25 assigning and occasionally modifying a participation

- 162 -

priority between the first RF communication network and the second RF communication network.

1/53




Figure 1a

SUBSTITUTE SHEET (RULE 26)

2/53




Figure 1b

3/53




Figure 1c




Figure 2

5/53

10 ms	10 ms	10 ms	10 ms	10 ms	10 ms	10 ms	10 ms
Freq. 1 2466 MHz	Freq. 2 2411 MHz	Freq. 3 2417 MHz	Freq. 4 2472 MHz	Freq. 5 2444 MHz	Freq. 79 2403 MHz	Freq. 1 2466 MHz
301	301	301	301	301	301	301	301

Figure 3a

20 ms	20 ms	20 ms	20 ms	20 ms	20 ms	20 ms	20 ms
Freq. 1 2466 MHz	Freq. 2 2411 MHz	Freq. 3 2417 MHz	Freq. 4 2472 MHz	Freq. 5 2444 MHz	Freq. 79 2403 MHz	Freq. 1 2466 MHz
303	303	303	303	303	303	303	303

Figure 3b

Freq 1: 2466 MHz		Freq 2: 2411 MHz		Freq 3: 2417 MHz		Freq 5: 2444 MHz	
Al. 1 20 ms	Al. 2 20 ms	Al. 1 20 ms	Al. 2 20 ms	Al. 1 20 ms	Al. 2 20 ms	Al. 1 20 ms	Al. 2 20 ms

Figure 4a

401

Freq 1: 2466 MHz				Freq 2: 2411 MHz			
Al. 1 20 ms	Al. 2 20 ms	Al. 3 20 ms	Al. 4 20 ms	Al. 1 20 ms	Al. 2 20 ms	Al. 3 20 ms	Al. 4 20 ms

Figure 4b

401




Figure 5a




Figure 5b




Figure 6a




Figure 6b

9/53




Figure 7a




Figure 7b

11/53




Figure 8

12/53




Figure 9b

Figure 9a




Figure 10

14/53




Figure 11

SUBSTITUTE SHEET (RULE 26)




Figure 12




Figure 13

17/53

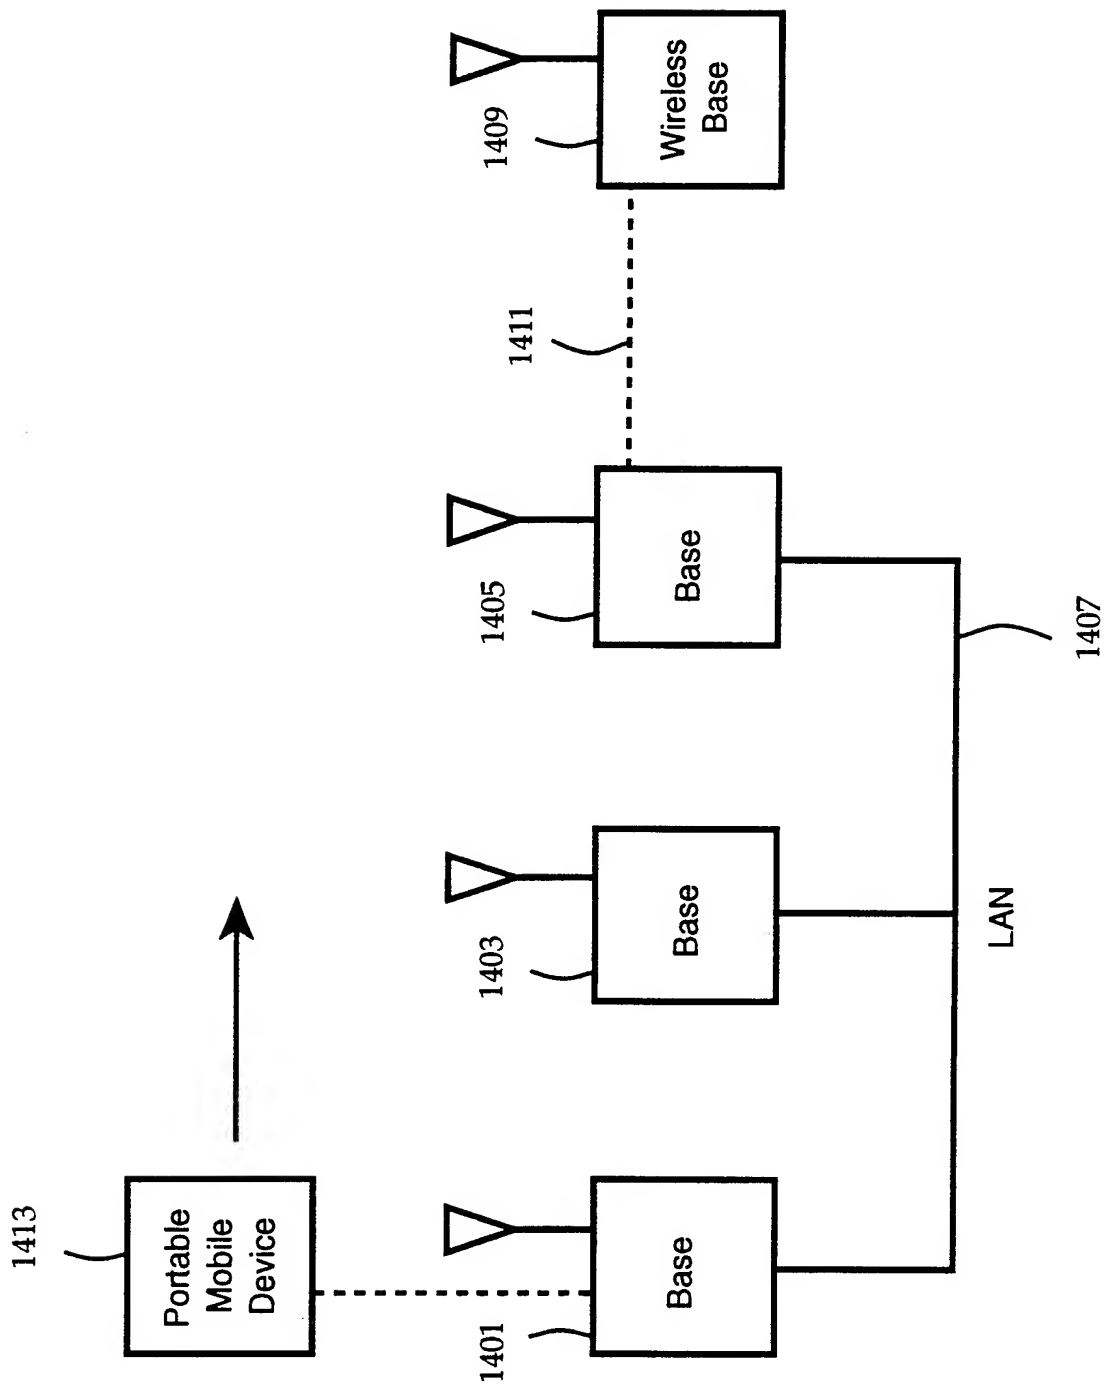


Figure 14




Figure 15

19/53

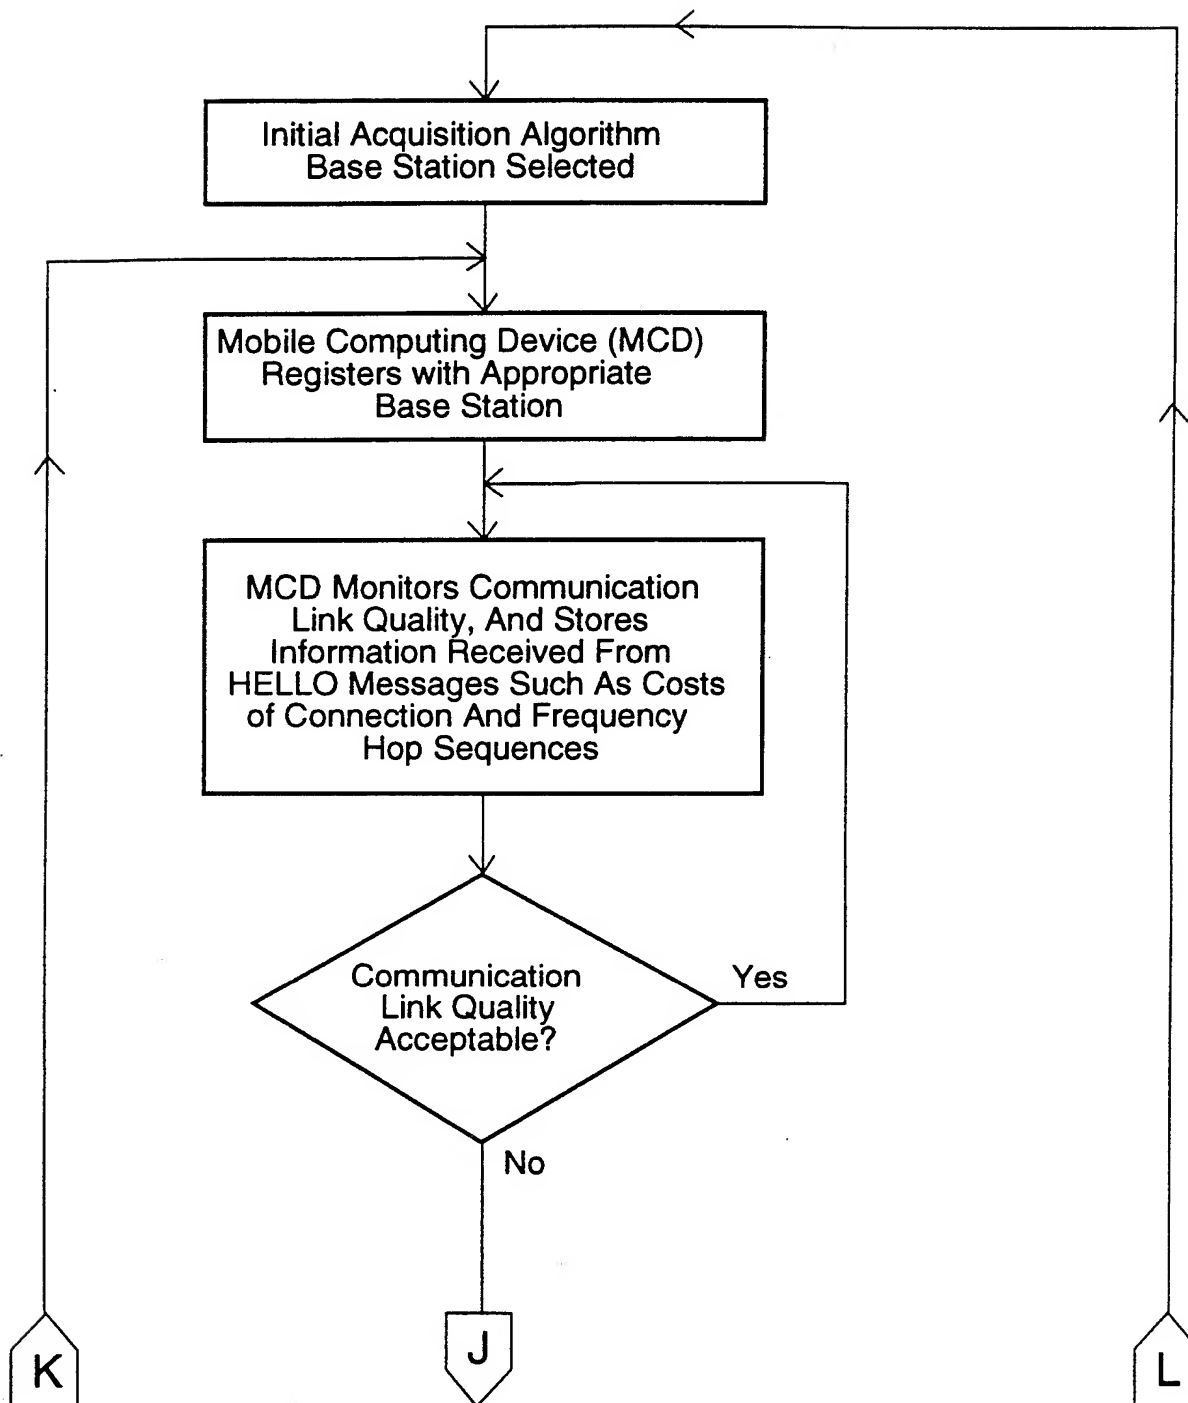


Figure 16a

SUBSTITUTE SHEET (RULE 26)

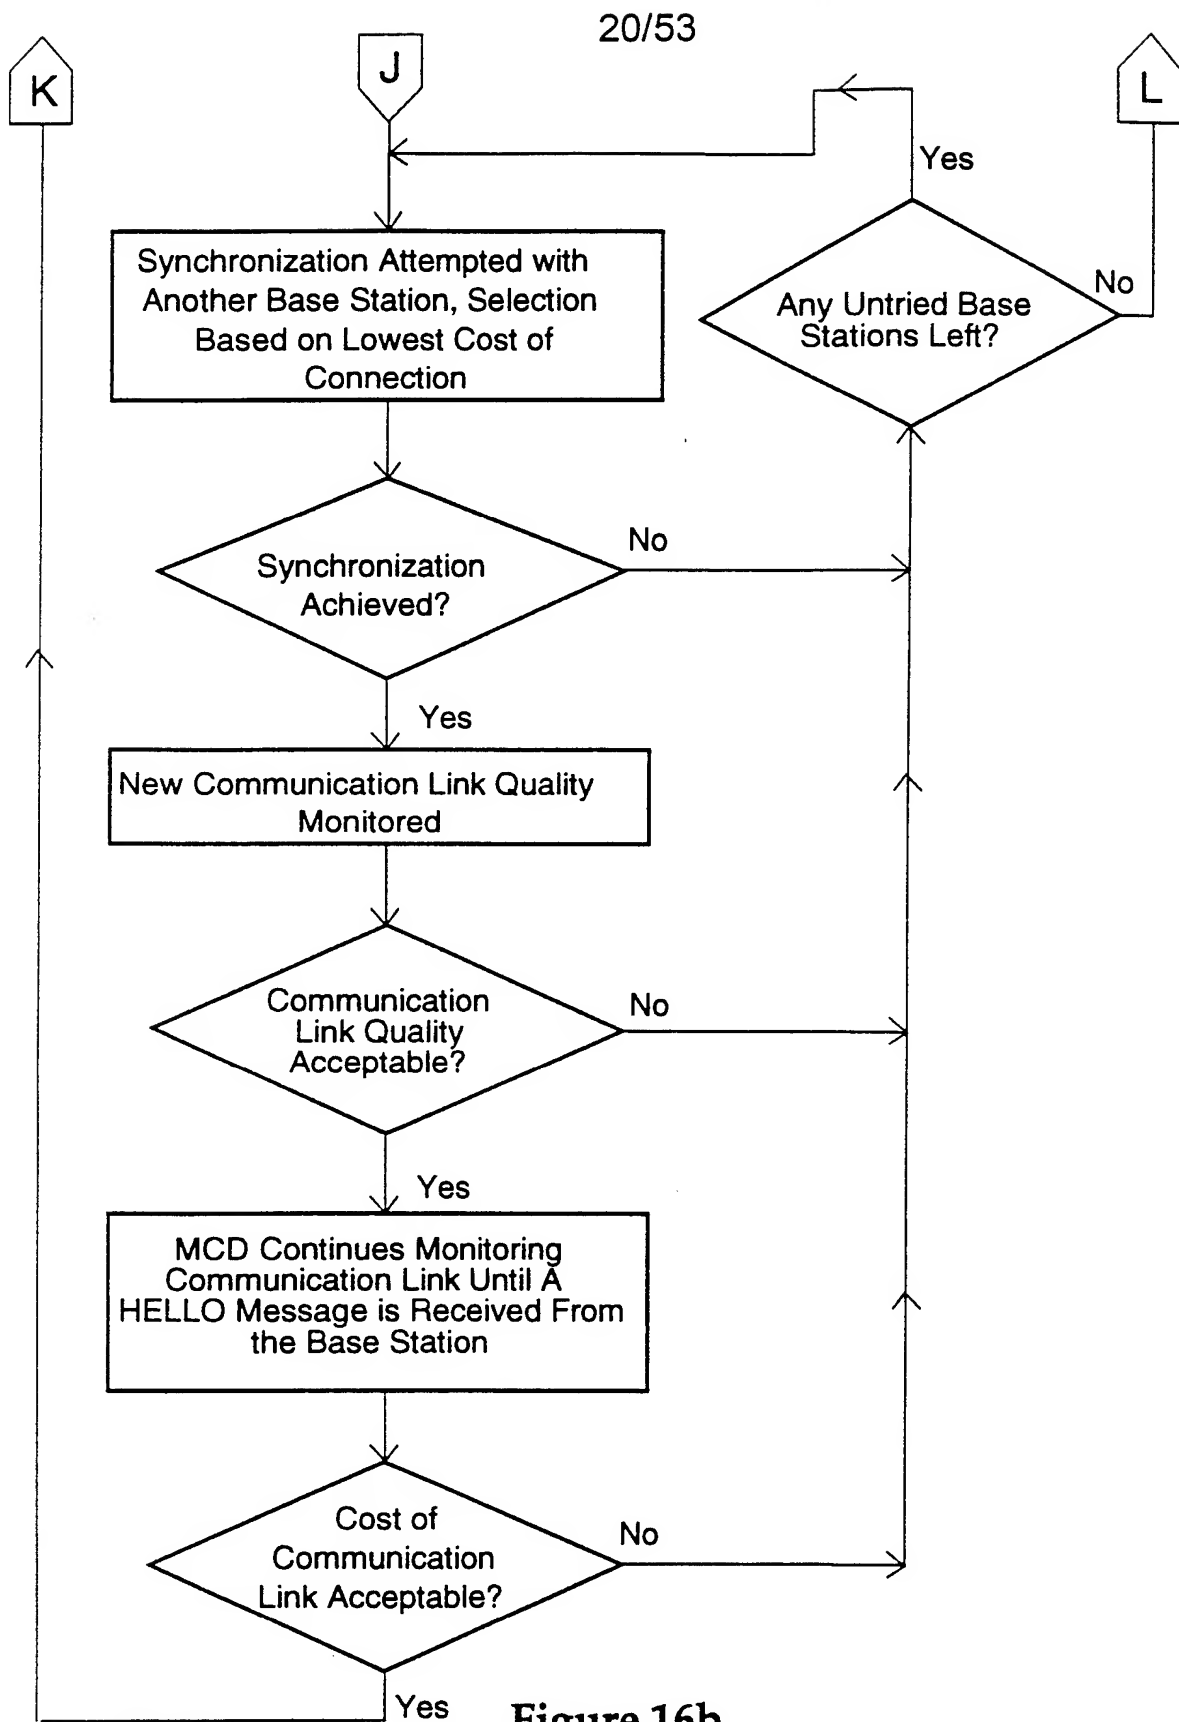


Figure 16b

21/53

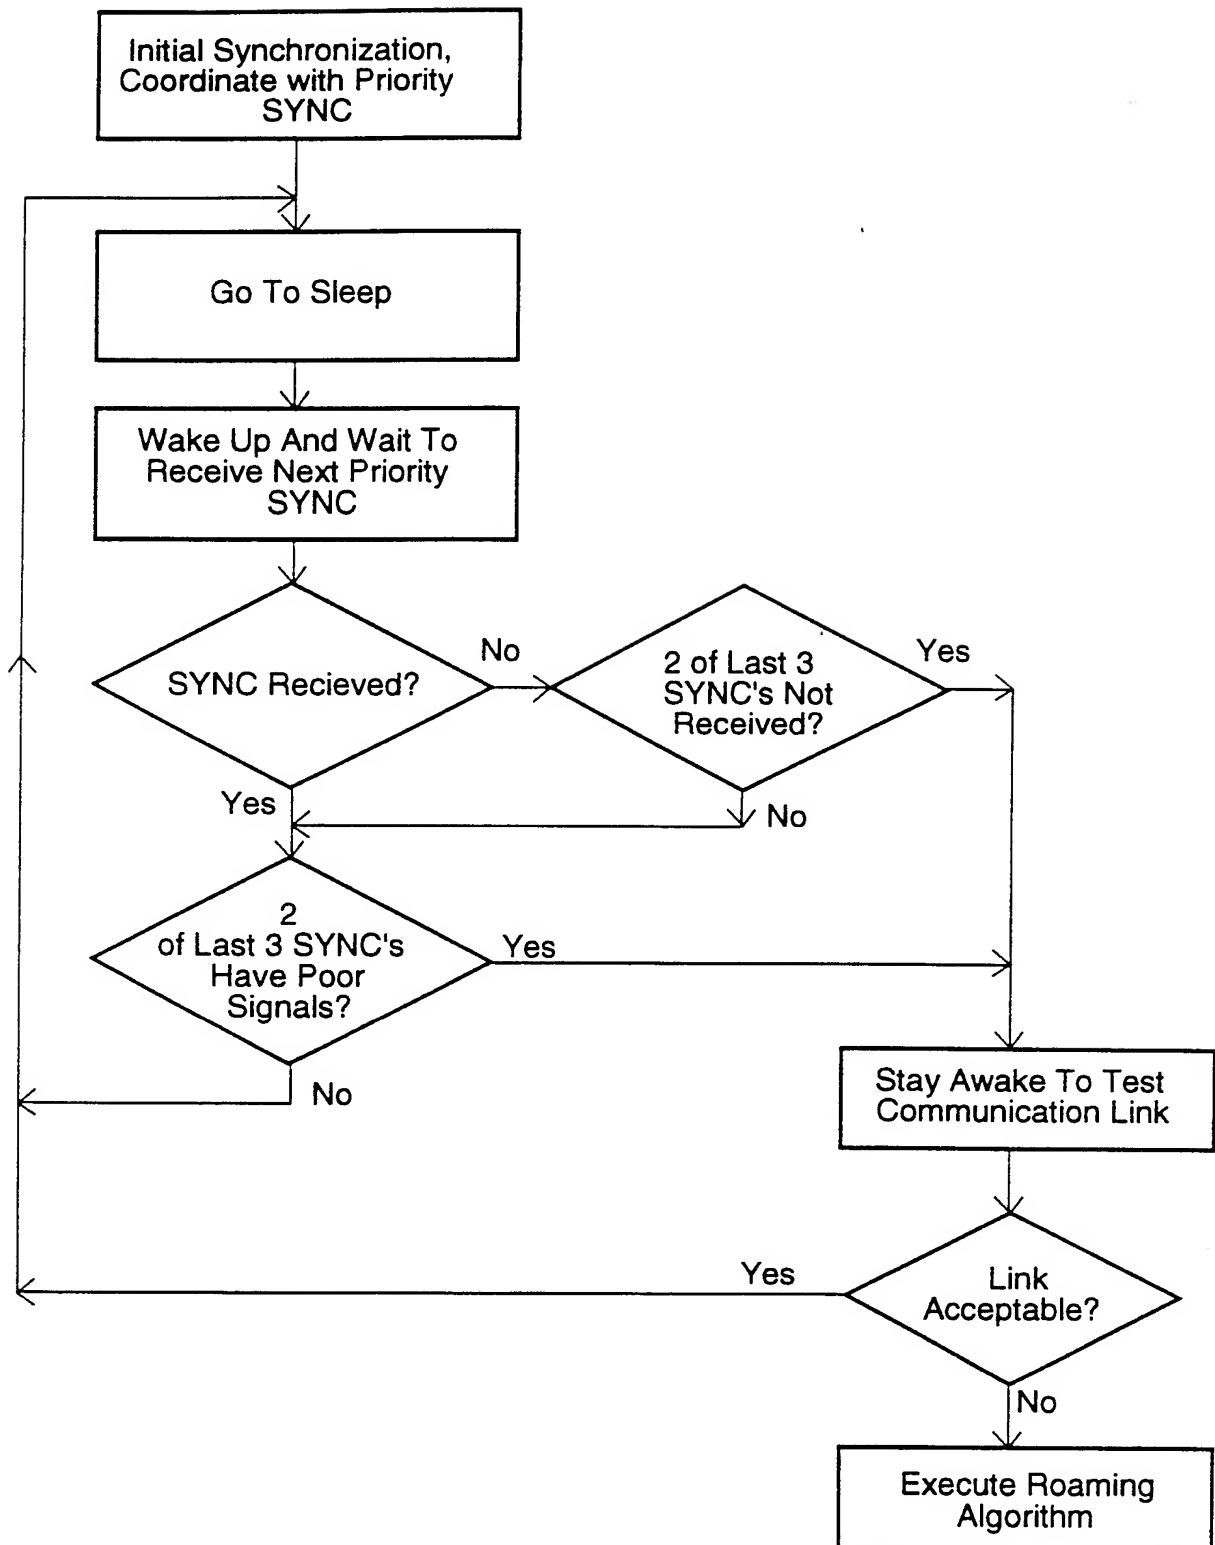


Figure 17

22/53

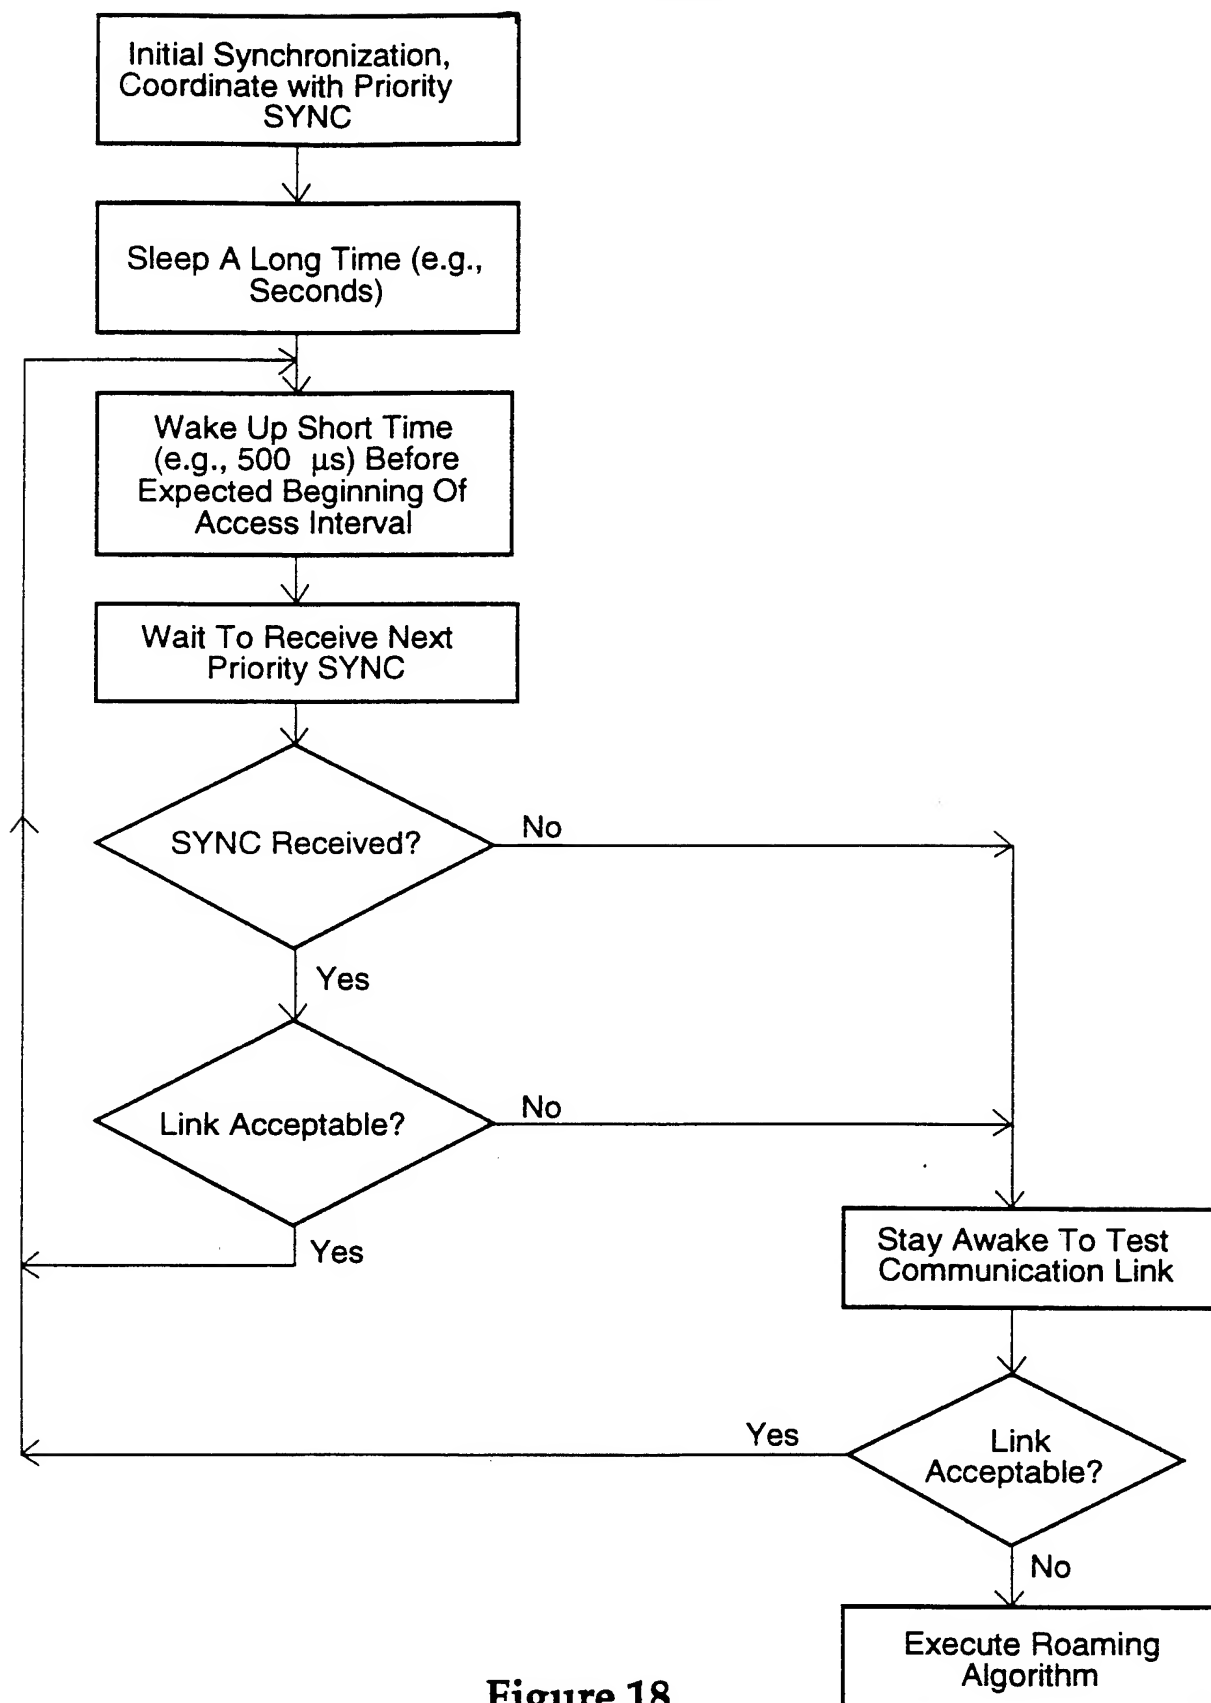


Figure 18

23/53

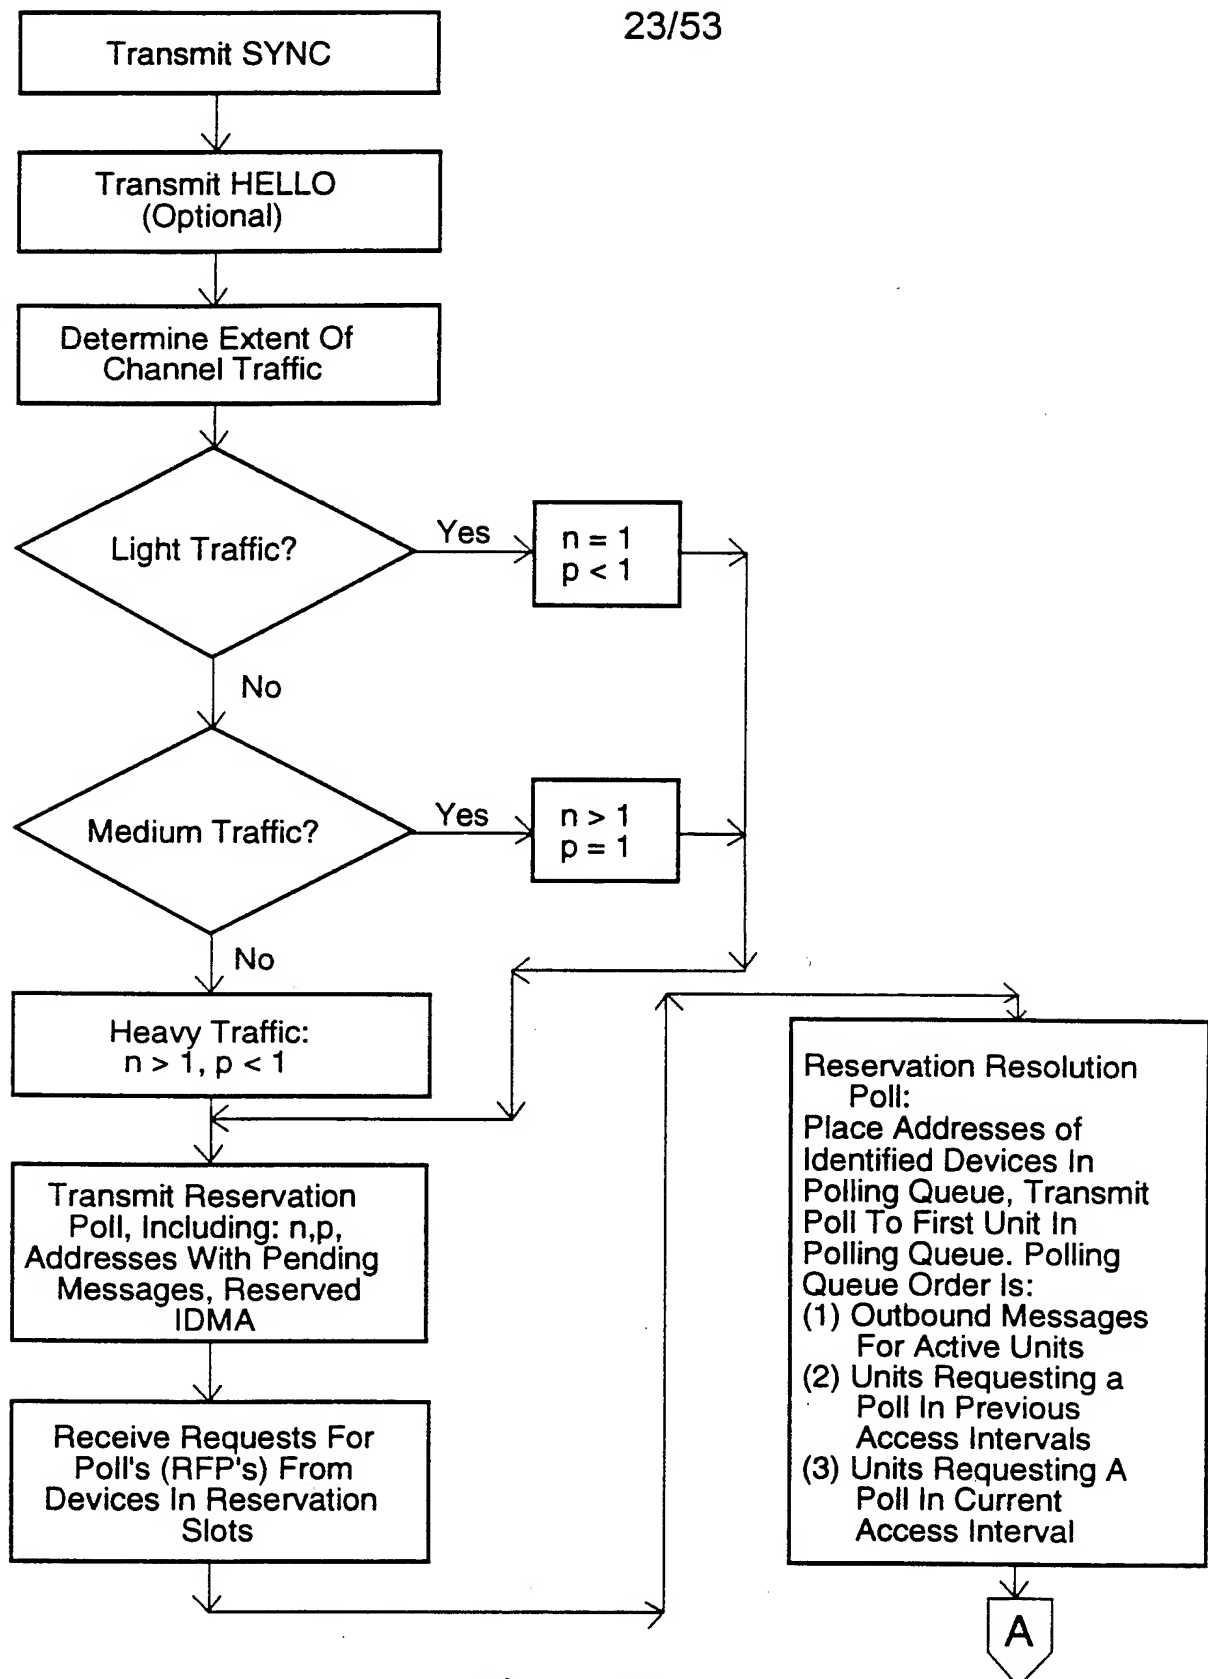


Figure 19a

24/53

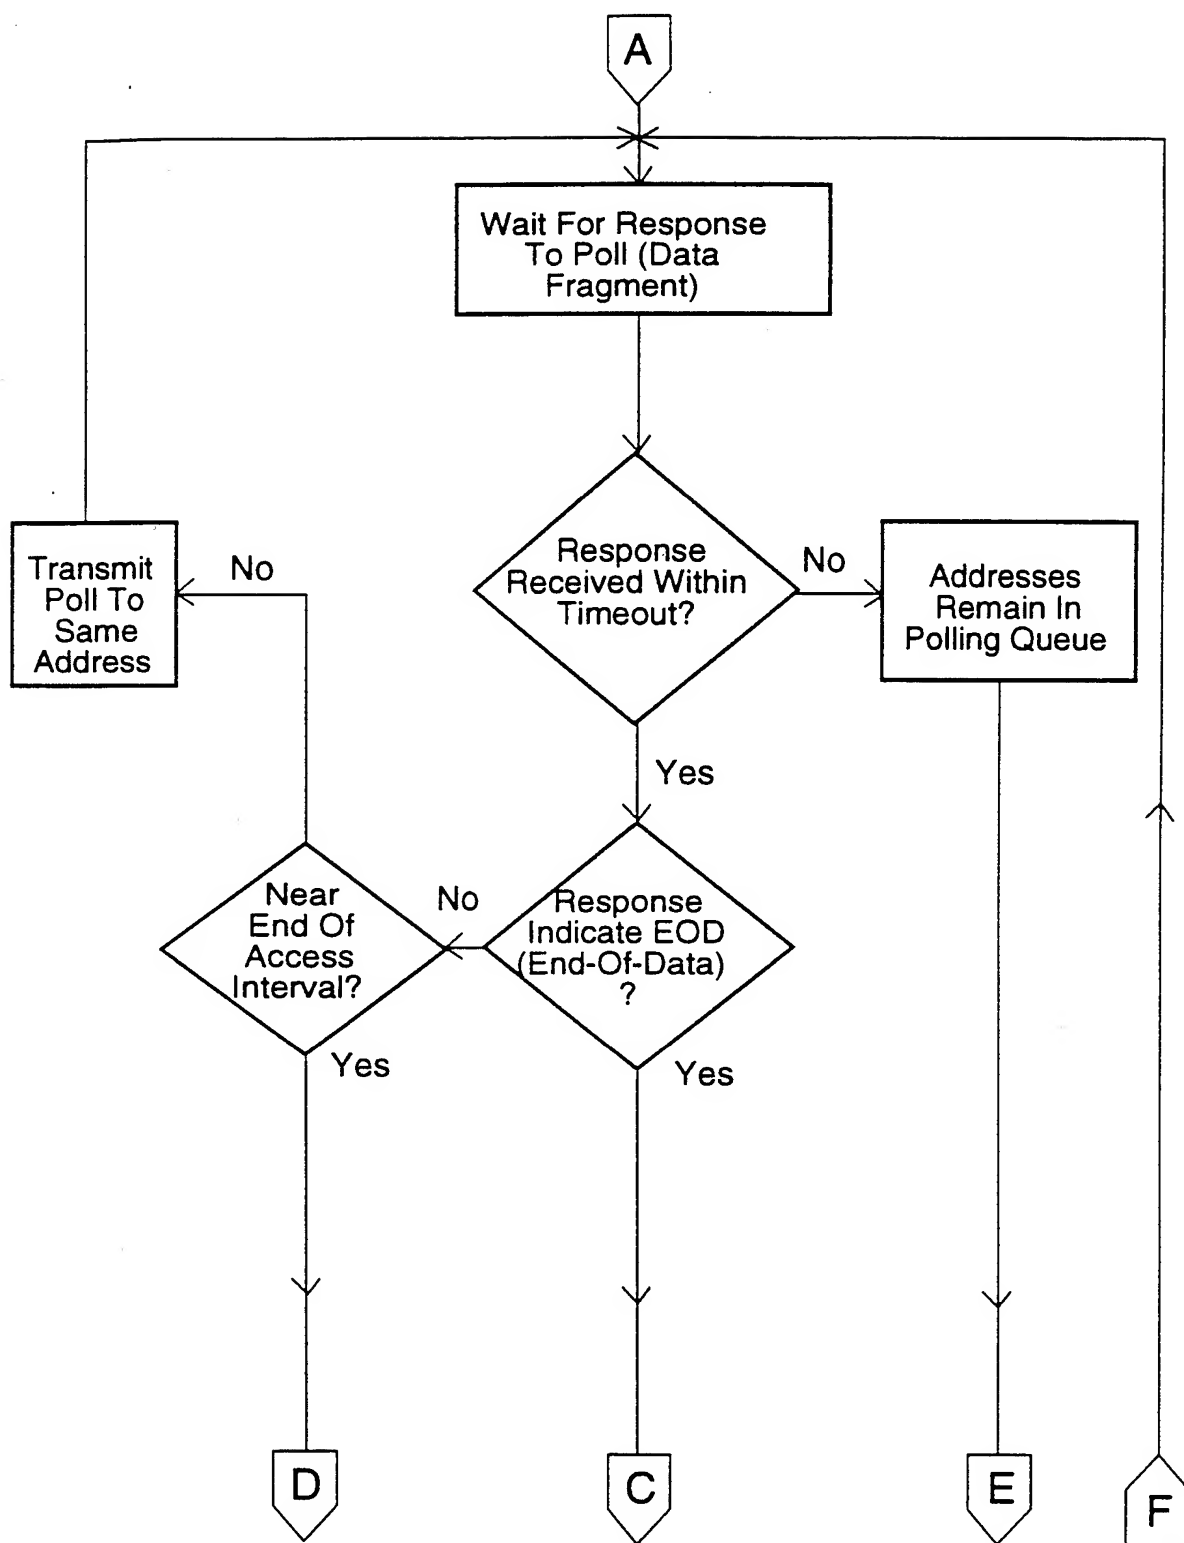


Figure 19b-1

25/53

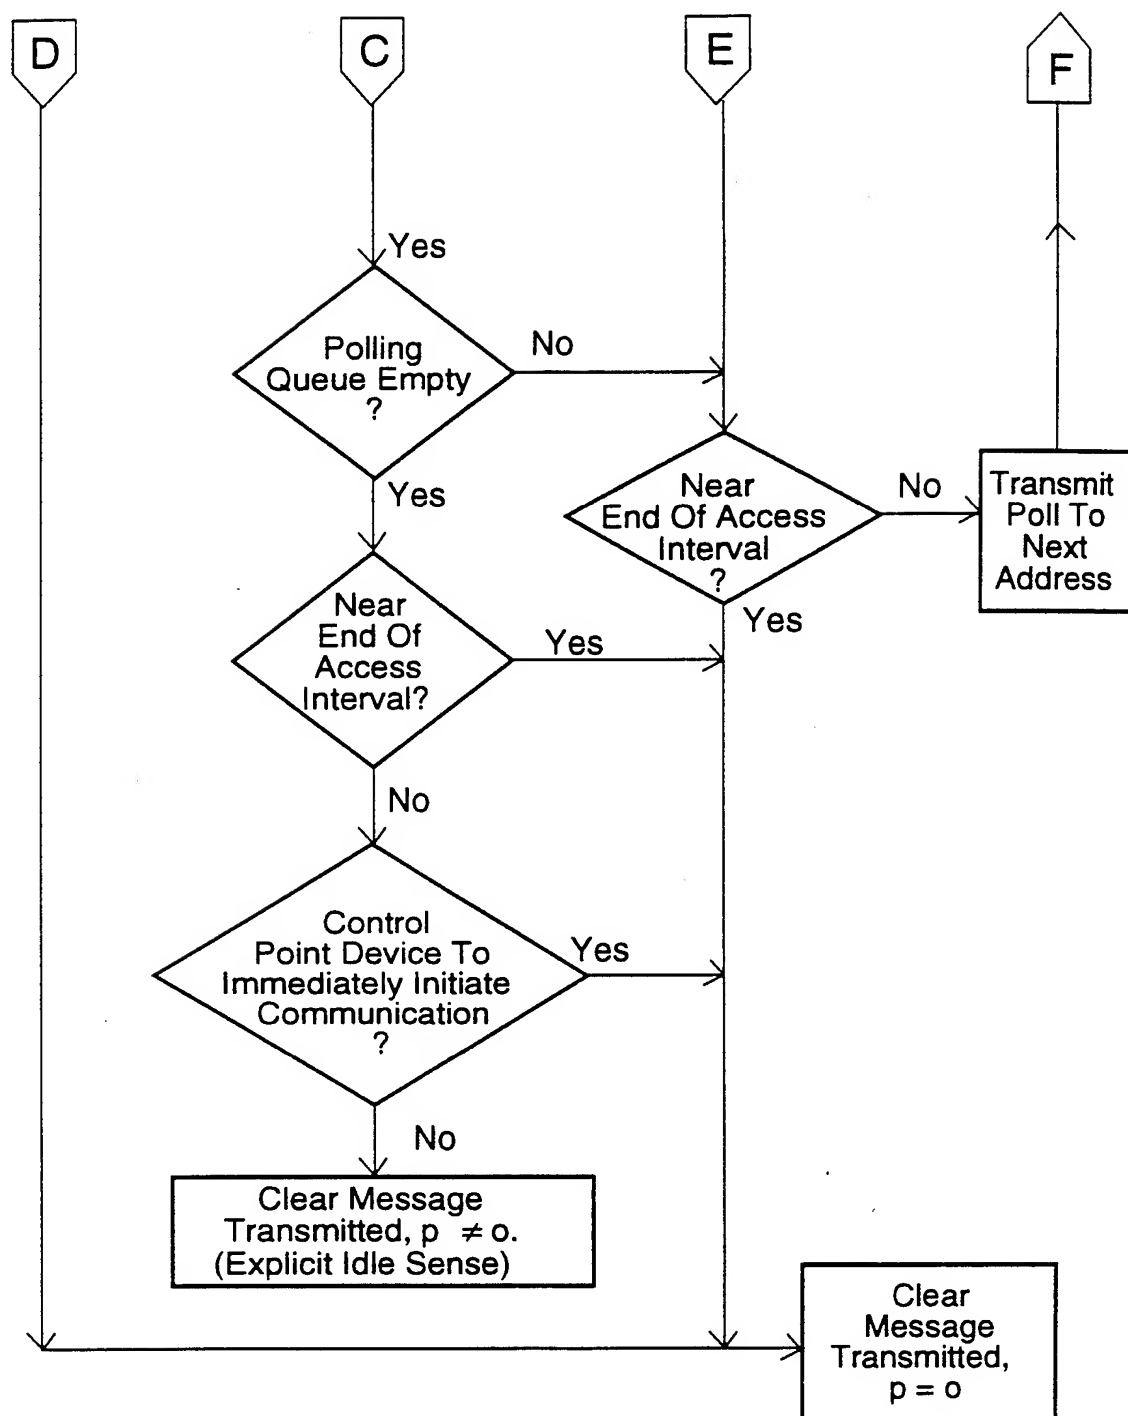


Figure 19b-2

26/53

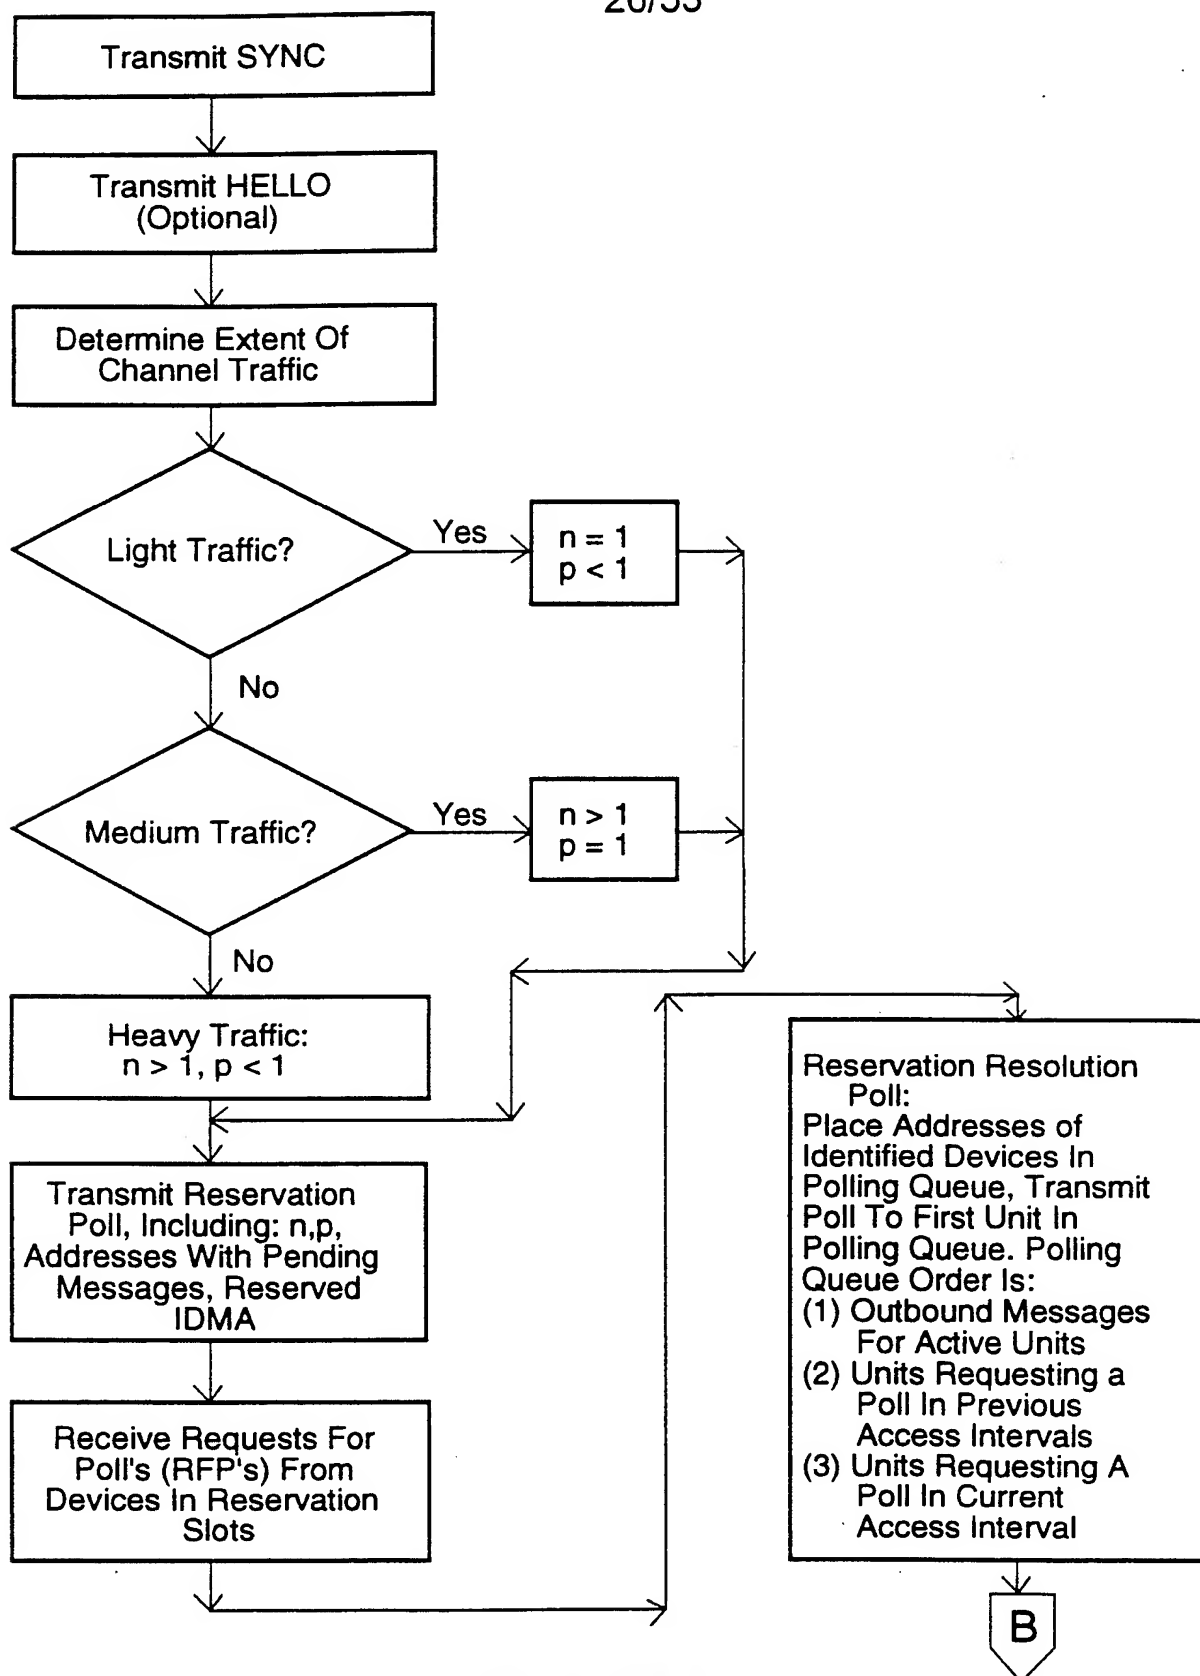


Figure 20a

27/53

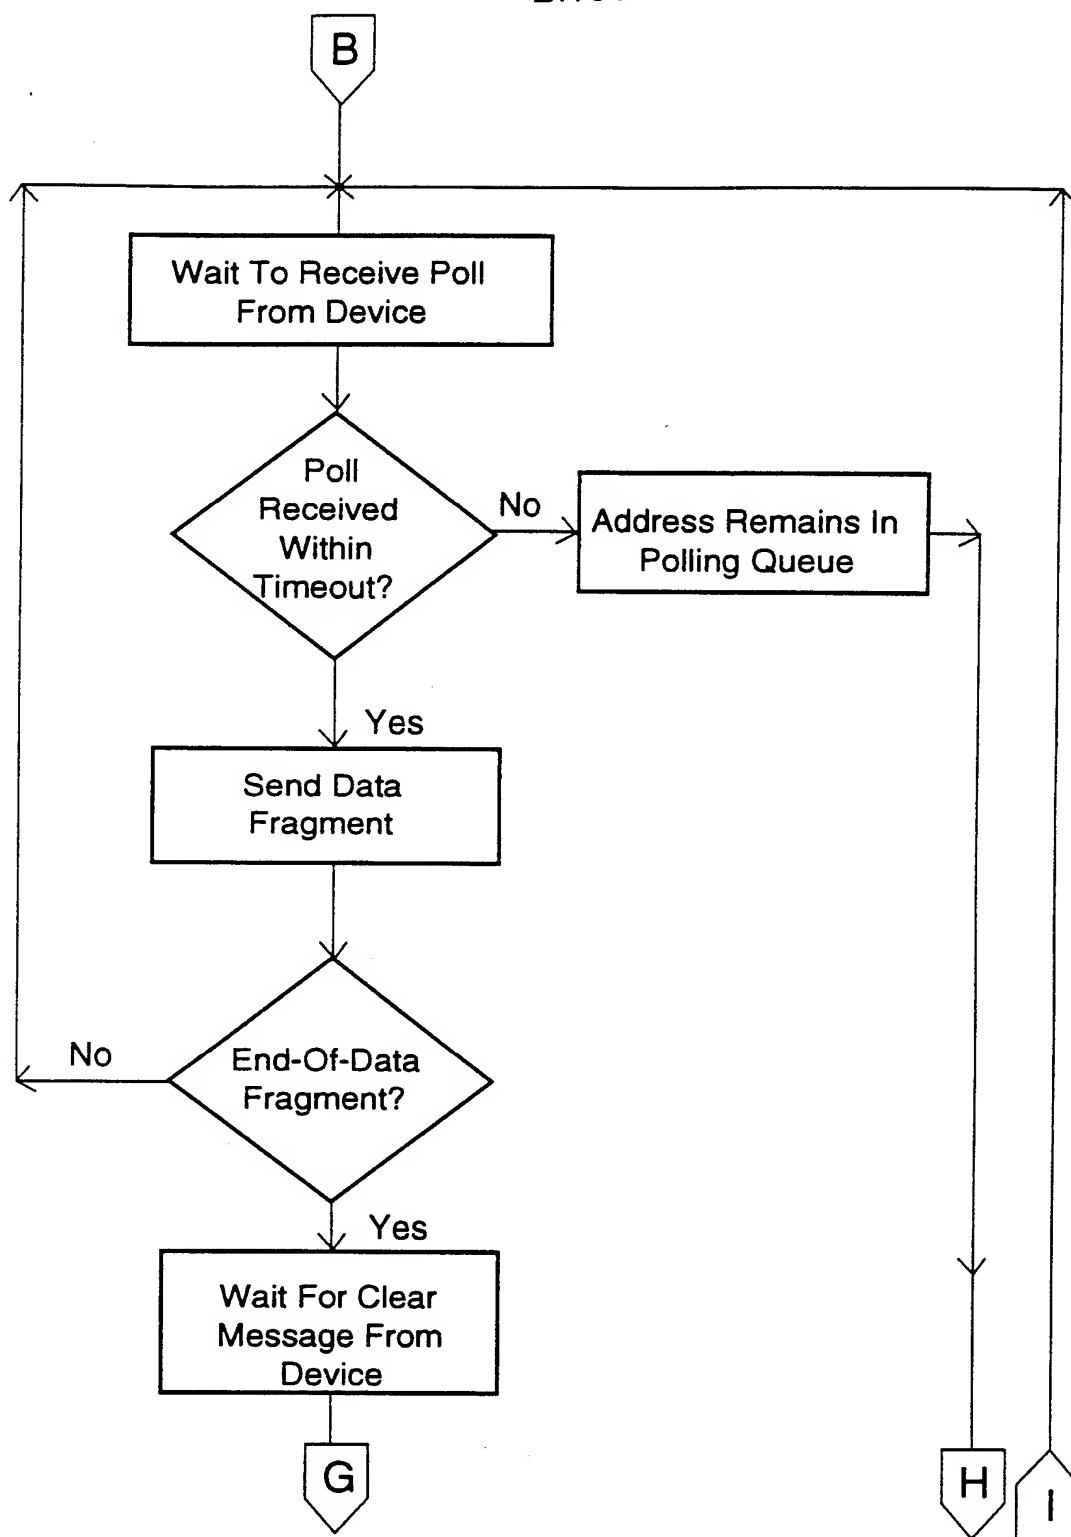


Figure 20b-1

SUBSTITUTE SHEET (RULE 26)

28/53

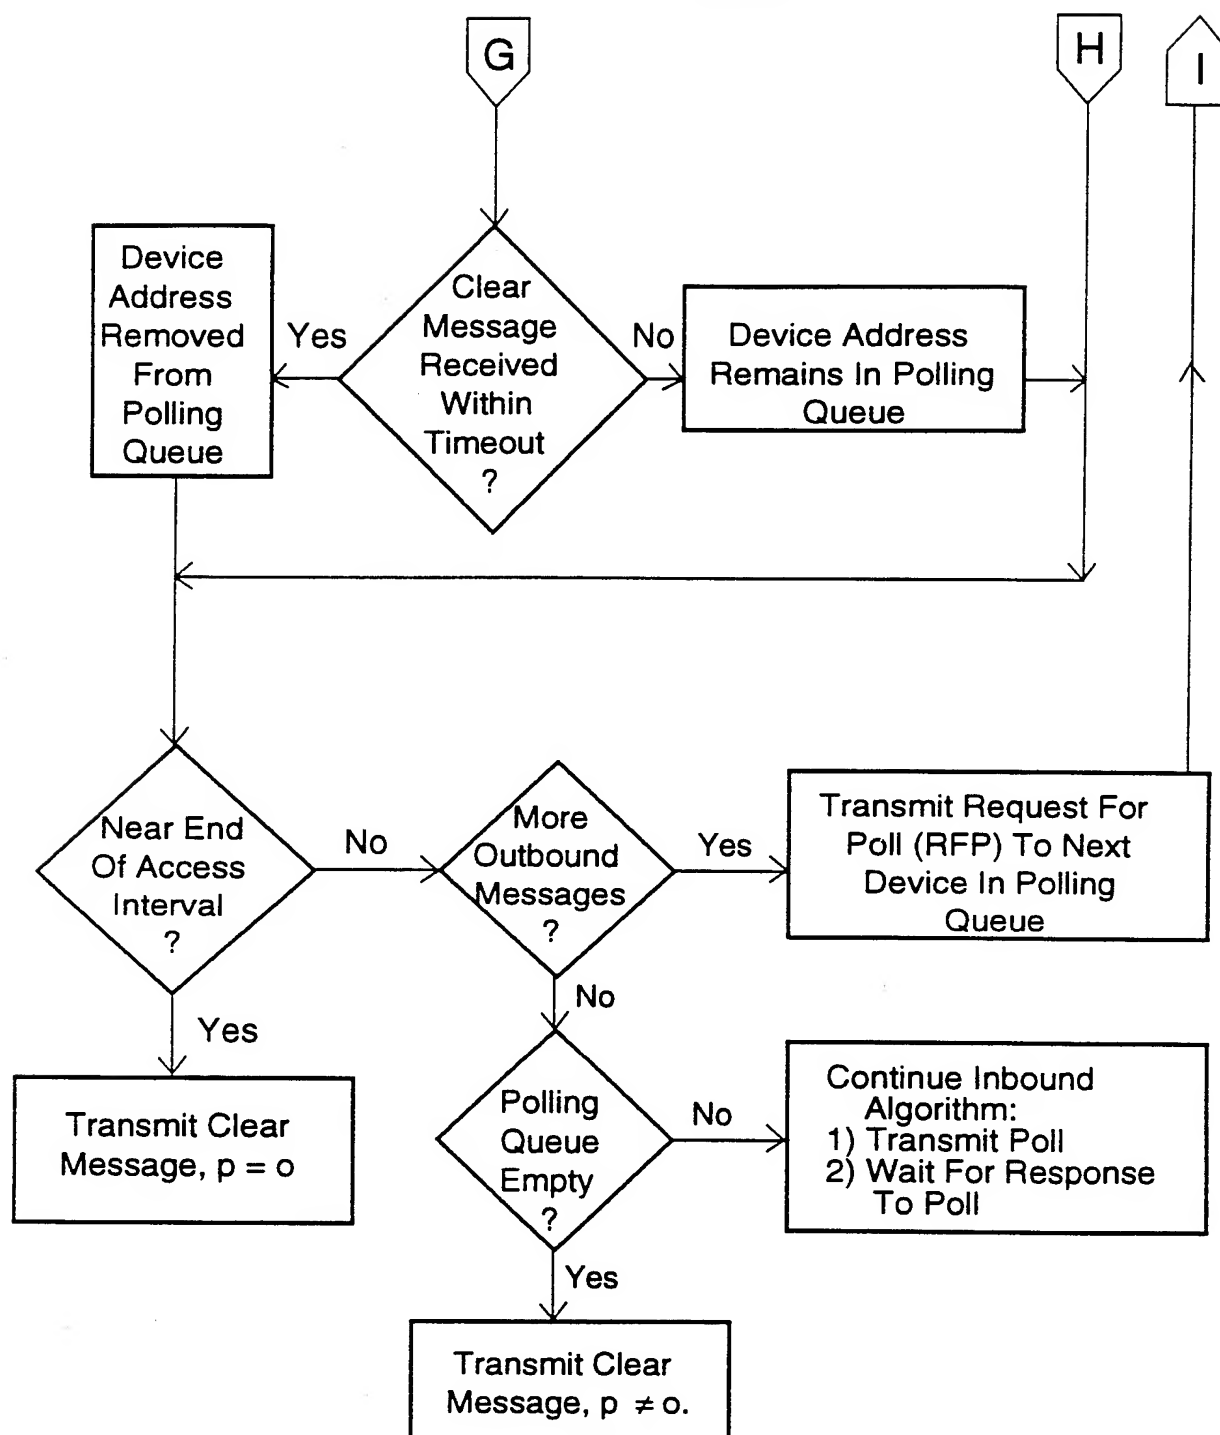


Figure 20b-2

29/53

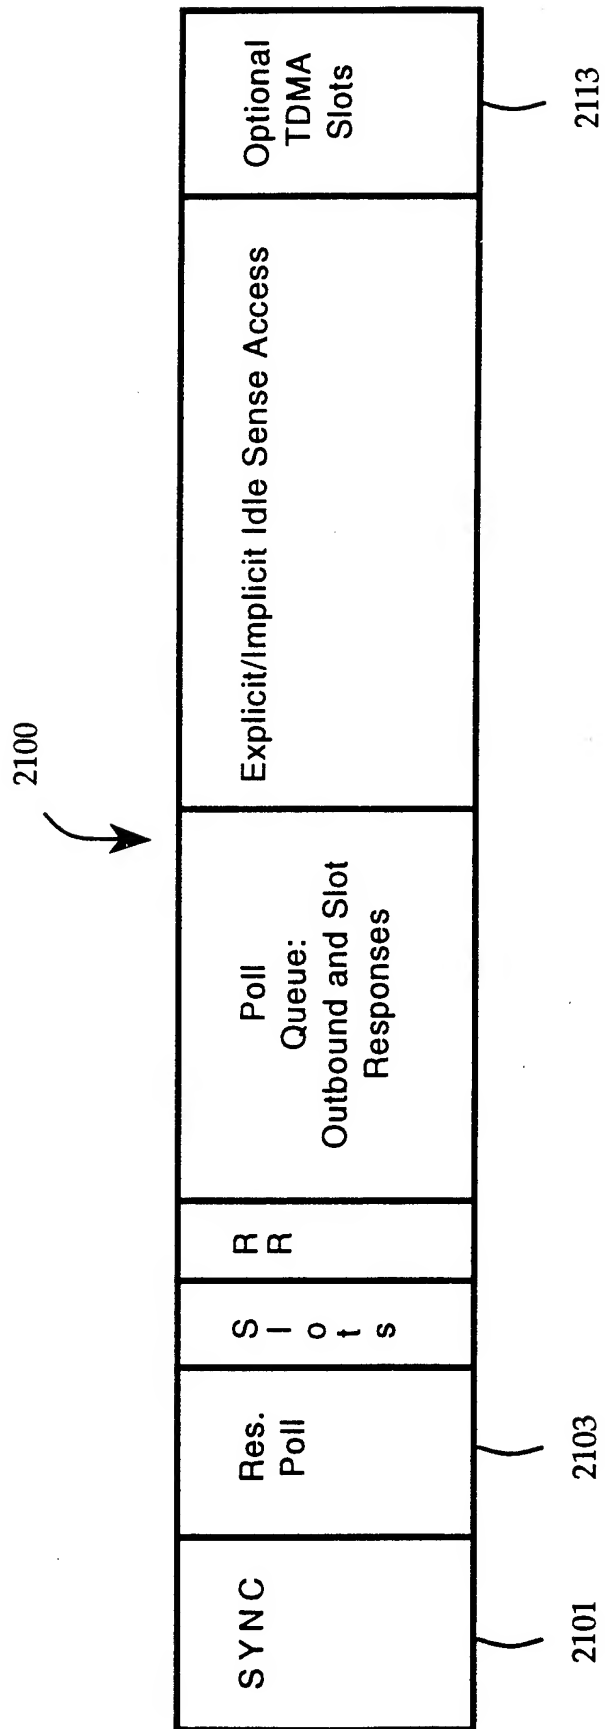


Figure 21

30/53




Figure 22




Figure 23

32/53




Figure 24




Figure 25

SYNC

2600

Byte	Content	Definition
1	FLAG	Frame Delimiter
2	Dest. Address 1	Destination Address (SYNC = Broadcast = FFFF _H)
3	Dest. Address 2	Destination Address, 2nd Byte
4	TYPE	Message Type=SYNC, Priority Sync, or Priority SYNC with HELLO
5	Source Address 1	Control Point Address
6	Source Address 2	Control Point Address, 2nd Byte
7*	SYSTEM ID	System Identifier (LAN ID)
8*	NET ID	Net ID
9*	Command/Status 1	Control Point Status
		2 Bits Control Point Type
		2 Bits Loading Factor
		1 Bit Scheduled Service Availability
		3 Bits Reserved
10*	Command/Status 2	Scheduled Services
		5 Bits Number of Allocated TDMA Time Slots
		3 Bits Number of Available Time Slots
11*	Random	Randomly Selected Character for Improved Encryption
12	CRC 1	Error Check
13	CRC 2	Error Check
14	FLAG	Frame Delimiter

Figure 26 *Optionally Encrypted Bytes

35/53

Reservation Poll

2700



Byte	Content	Definition
1	FLAG	Frame Delimiter
2	Dest. Address 1	Destination Address (RES = Broadcast = FFFF _H)
3	Dest. Address 2	Destination Address, 2nd Byte
4	TYPE	Message Type = RES POLL
5	SLOT PARMS	3 Bits Number of Slots
6		5 Bits Probability Factor
7*	Sequence	Hopping Sequence ID
8*	Index	Hopping Sequence Index
9*	NET PARMS	
		3 Bits Hop Rate
		2 Bits AI's per Hop
		3 Bits Reserved
10	CRC 1	Error Check
11	CRC 2	Error Check
12	FLAG	Frame Delimiter

Figure 27




Fig. 28A

37/53




Fig. 28b

38/53




Figure 28C

39/53




Fig. 29

40/53




Fig. 30

41/53








Fig. 31

42/53

Fig. 32

43/53

**Fig. 33a****Fig. 33b****Fig. 33c**

44/53




Fig. 34

45/53

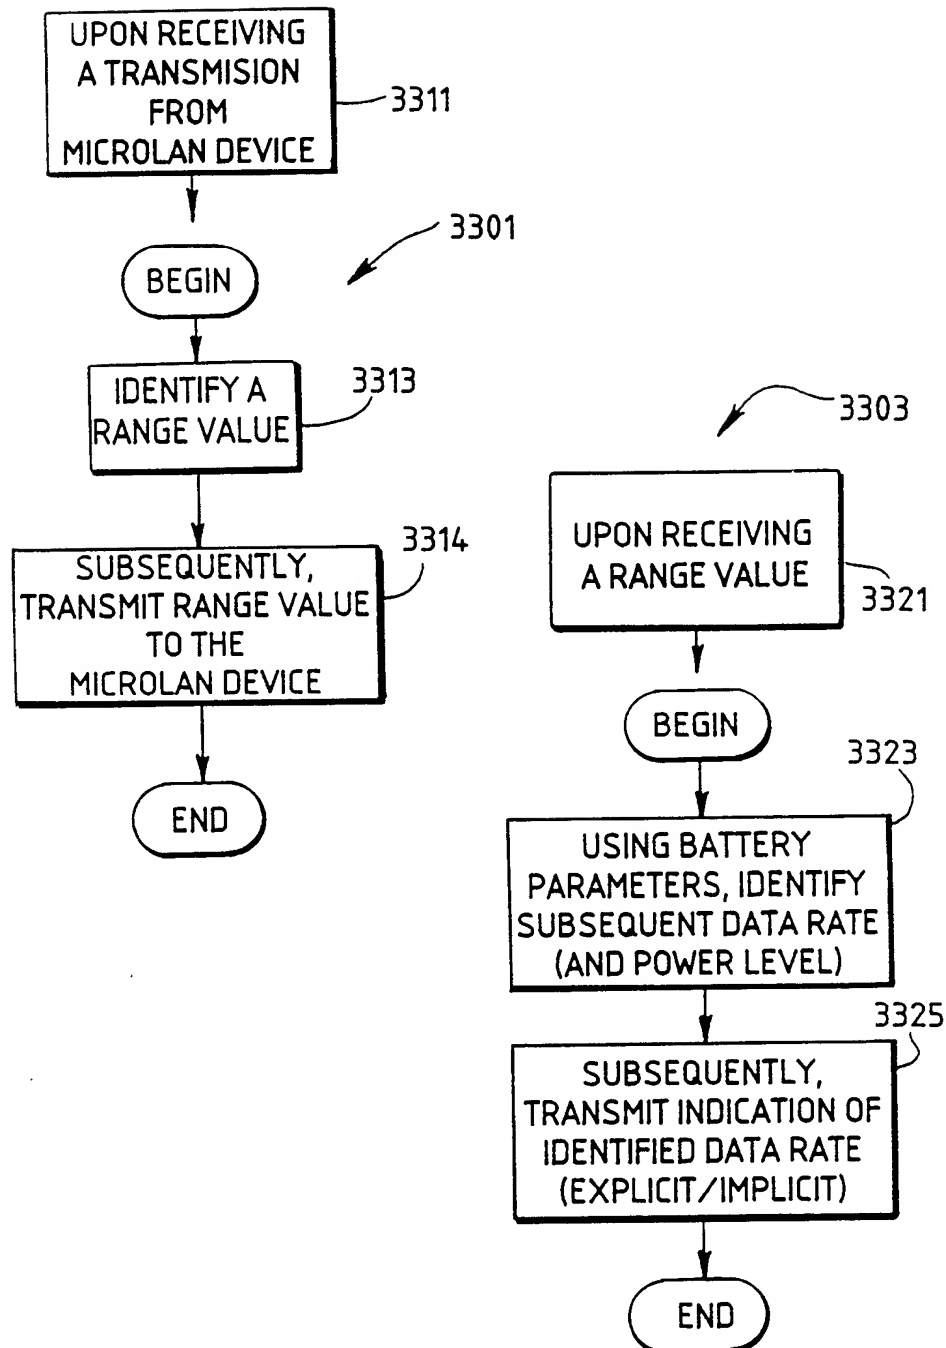
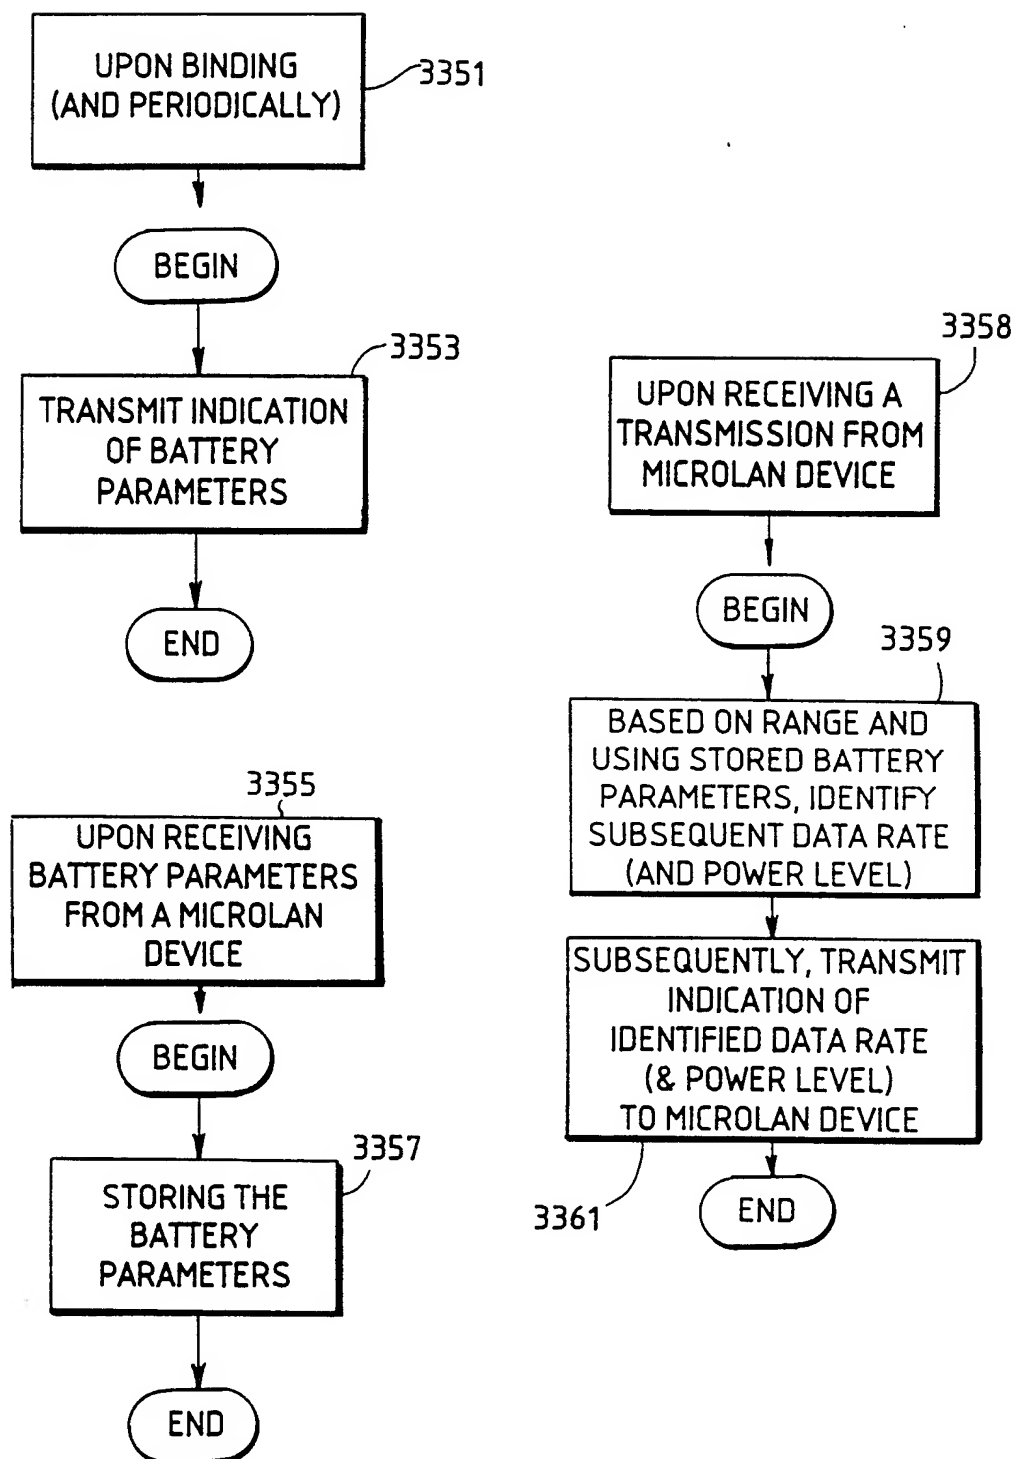


Fig. 35

46/53

Fig. 36

47/53




Figure 37

48/53




Figure 38

49/53




Figure 39




Figure 40

50/53

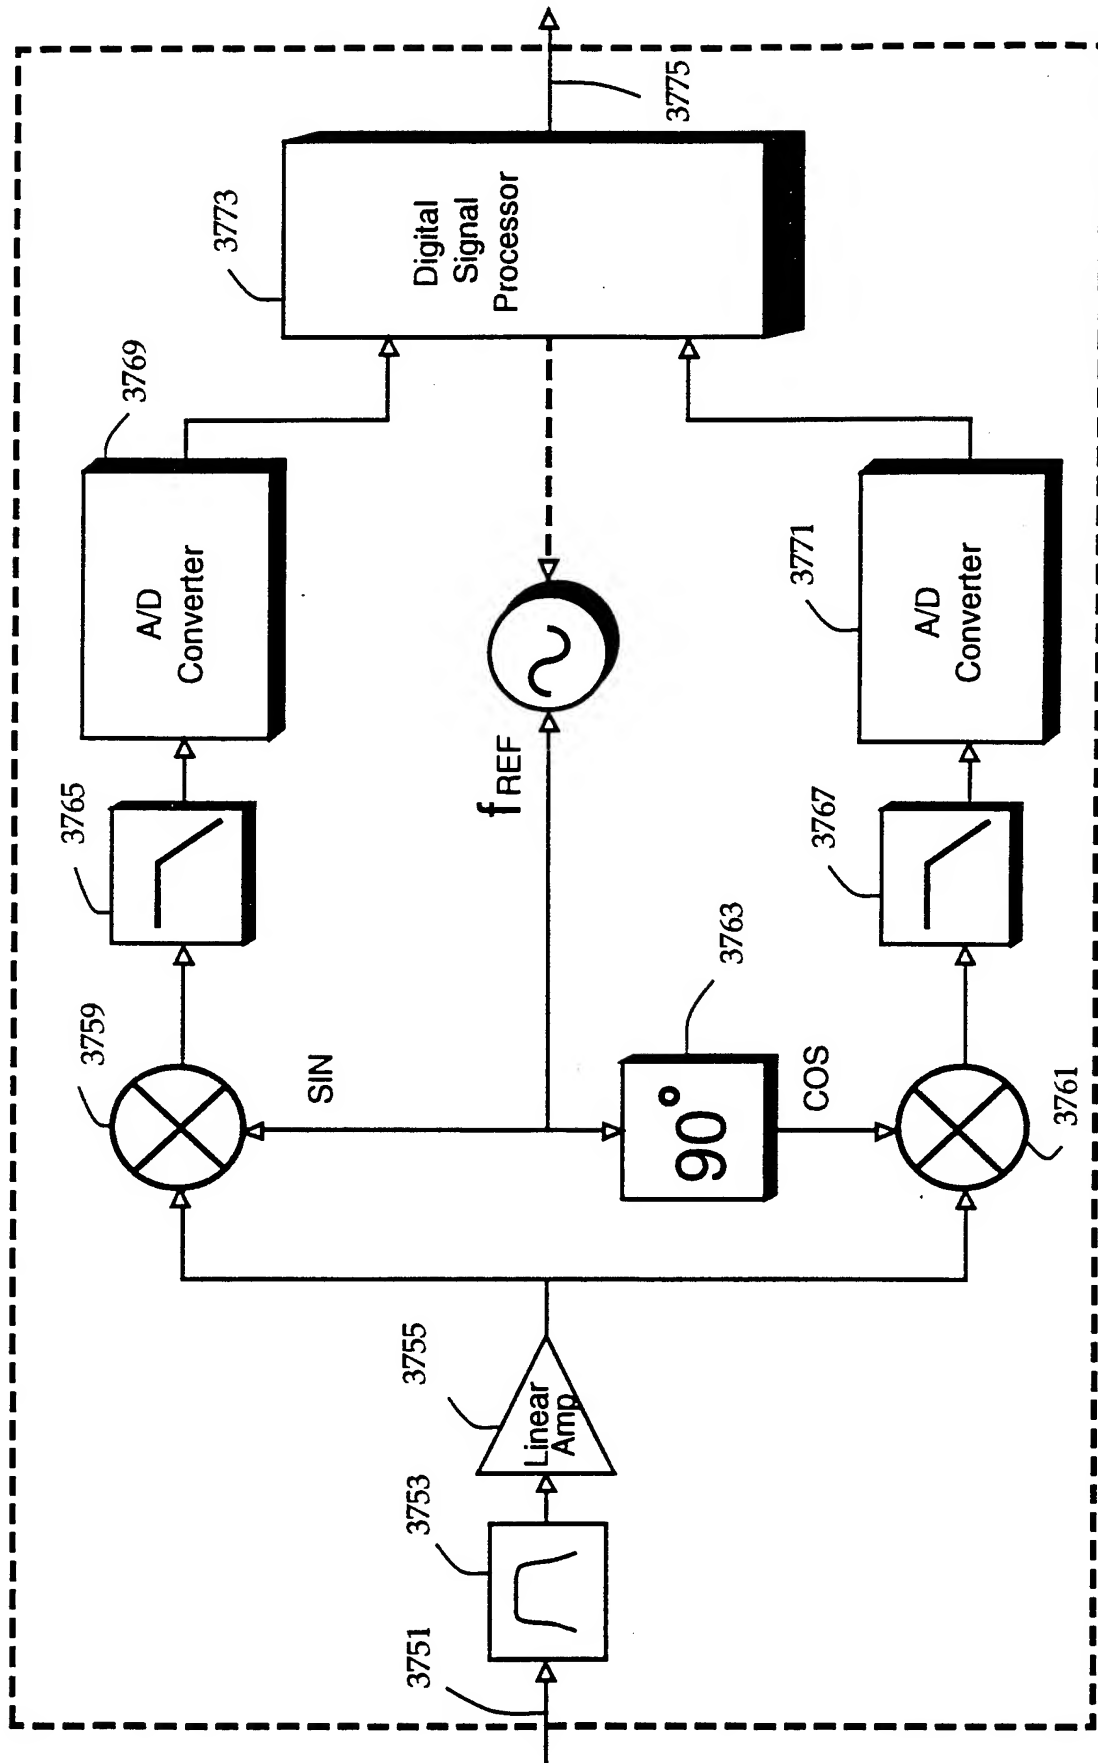
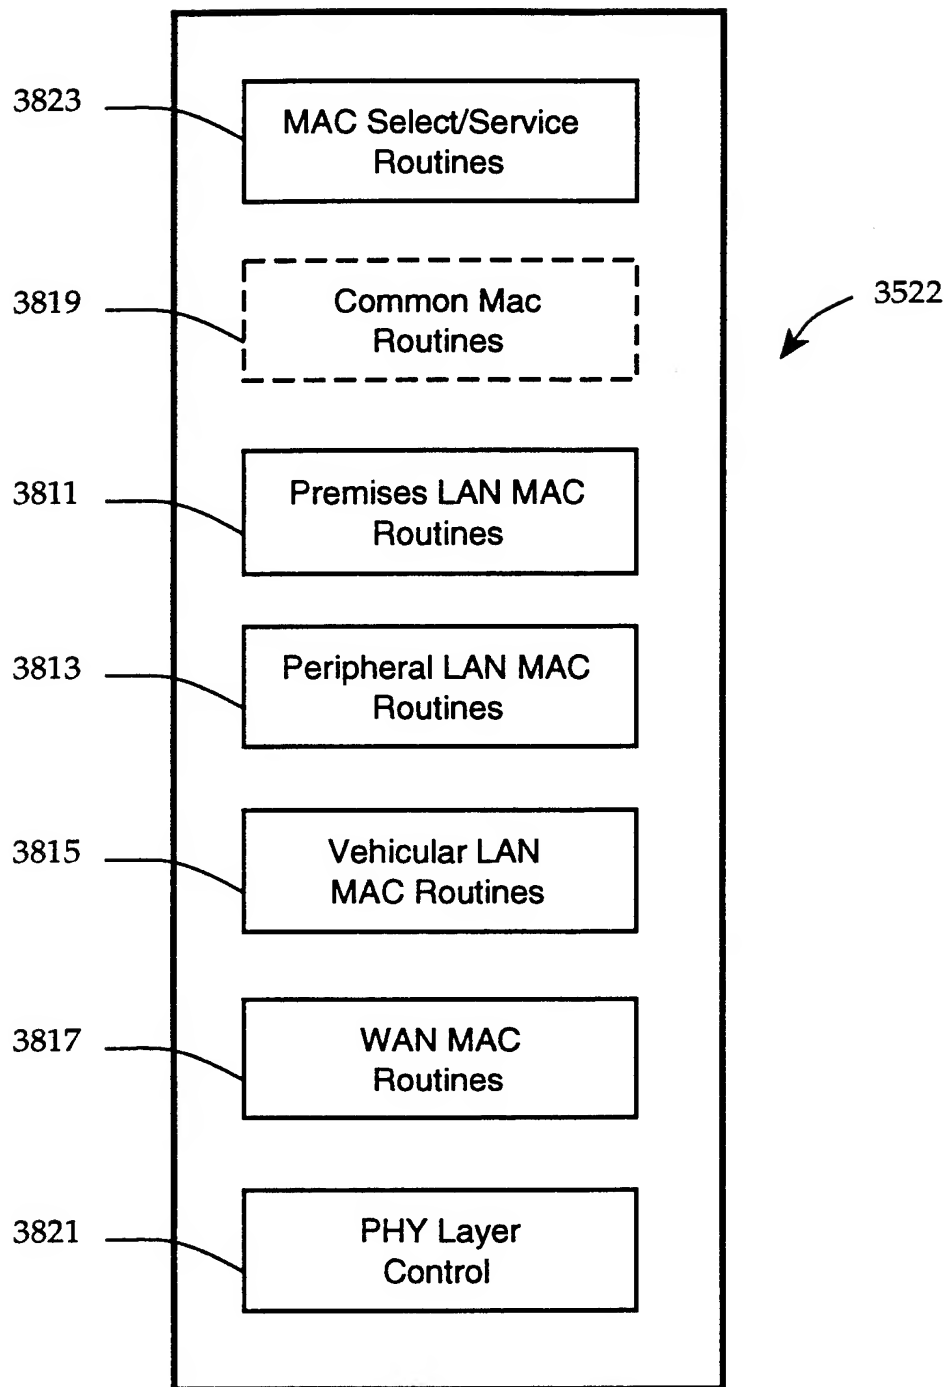


Figure 41

51/53

**Figure 42**

52/53

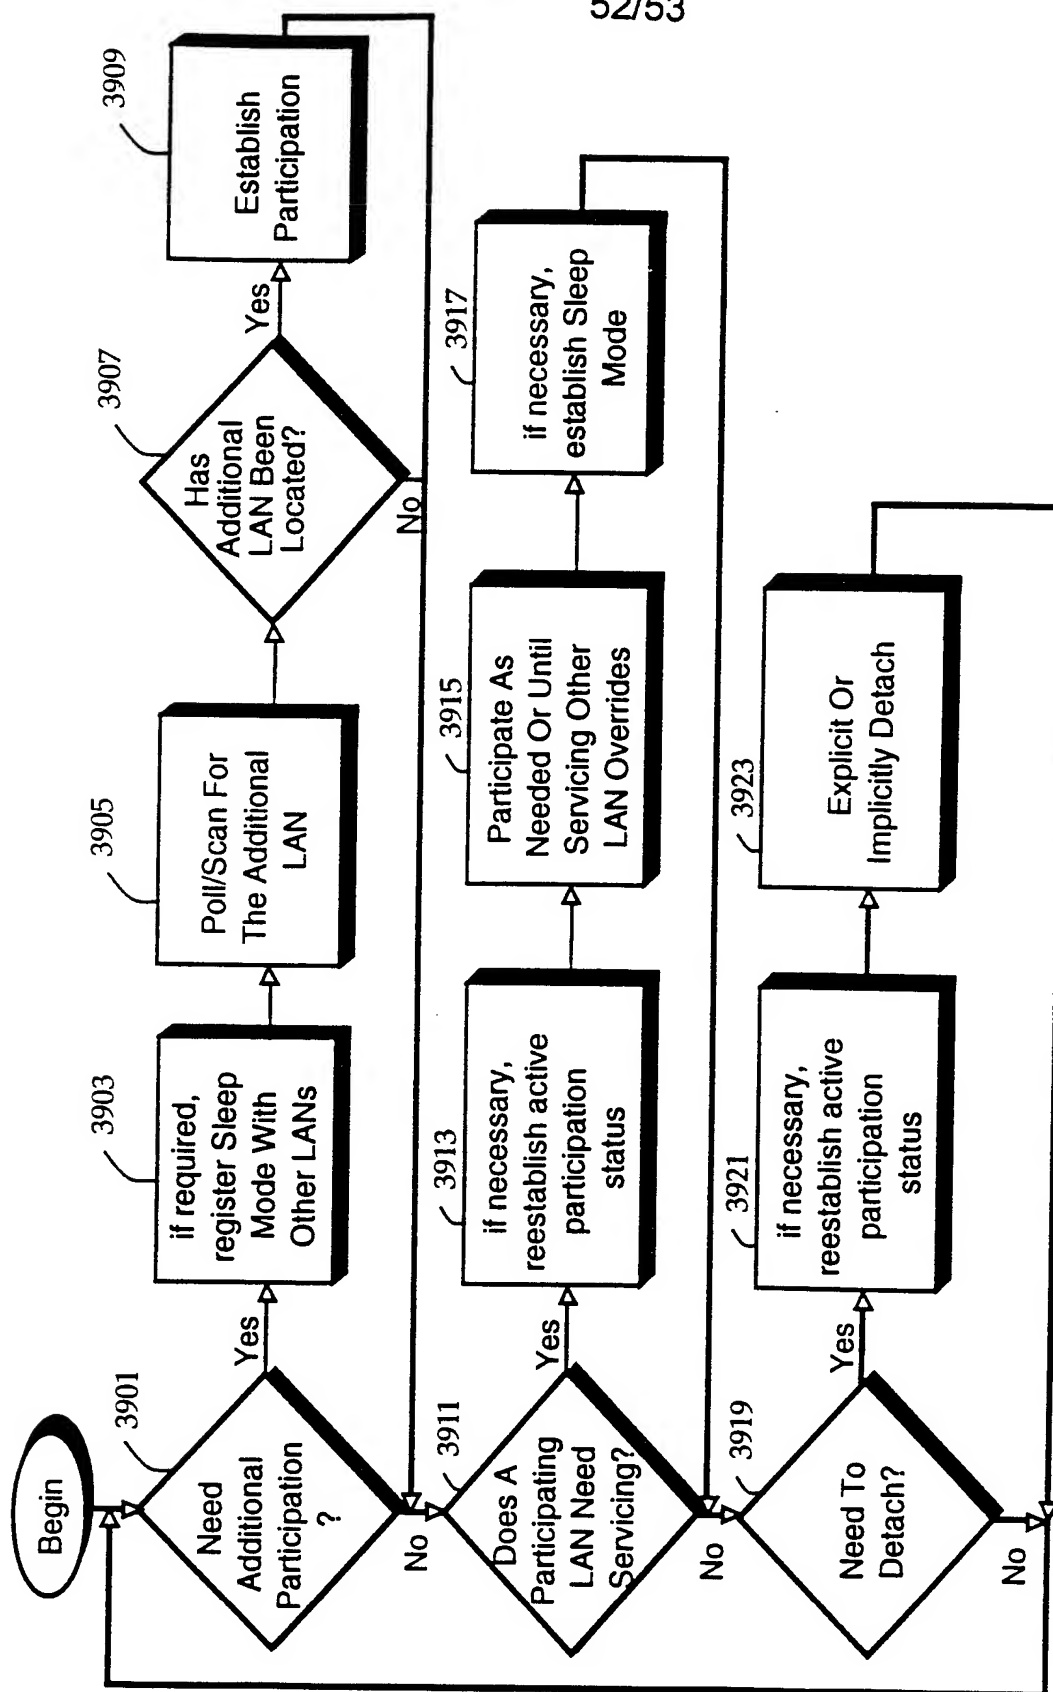


Figure 43

53/53

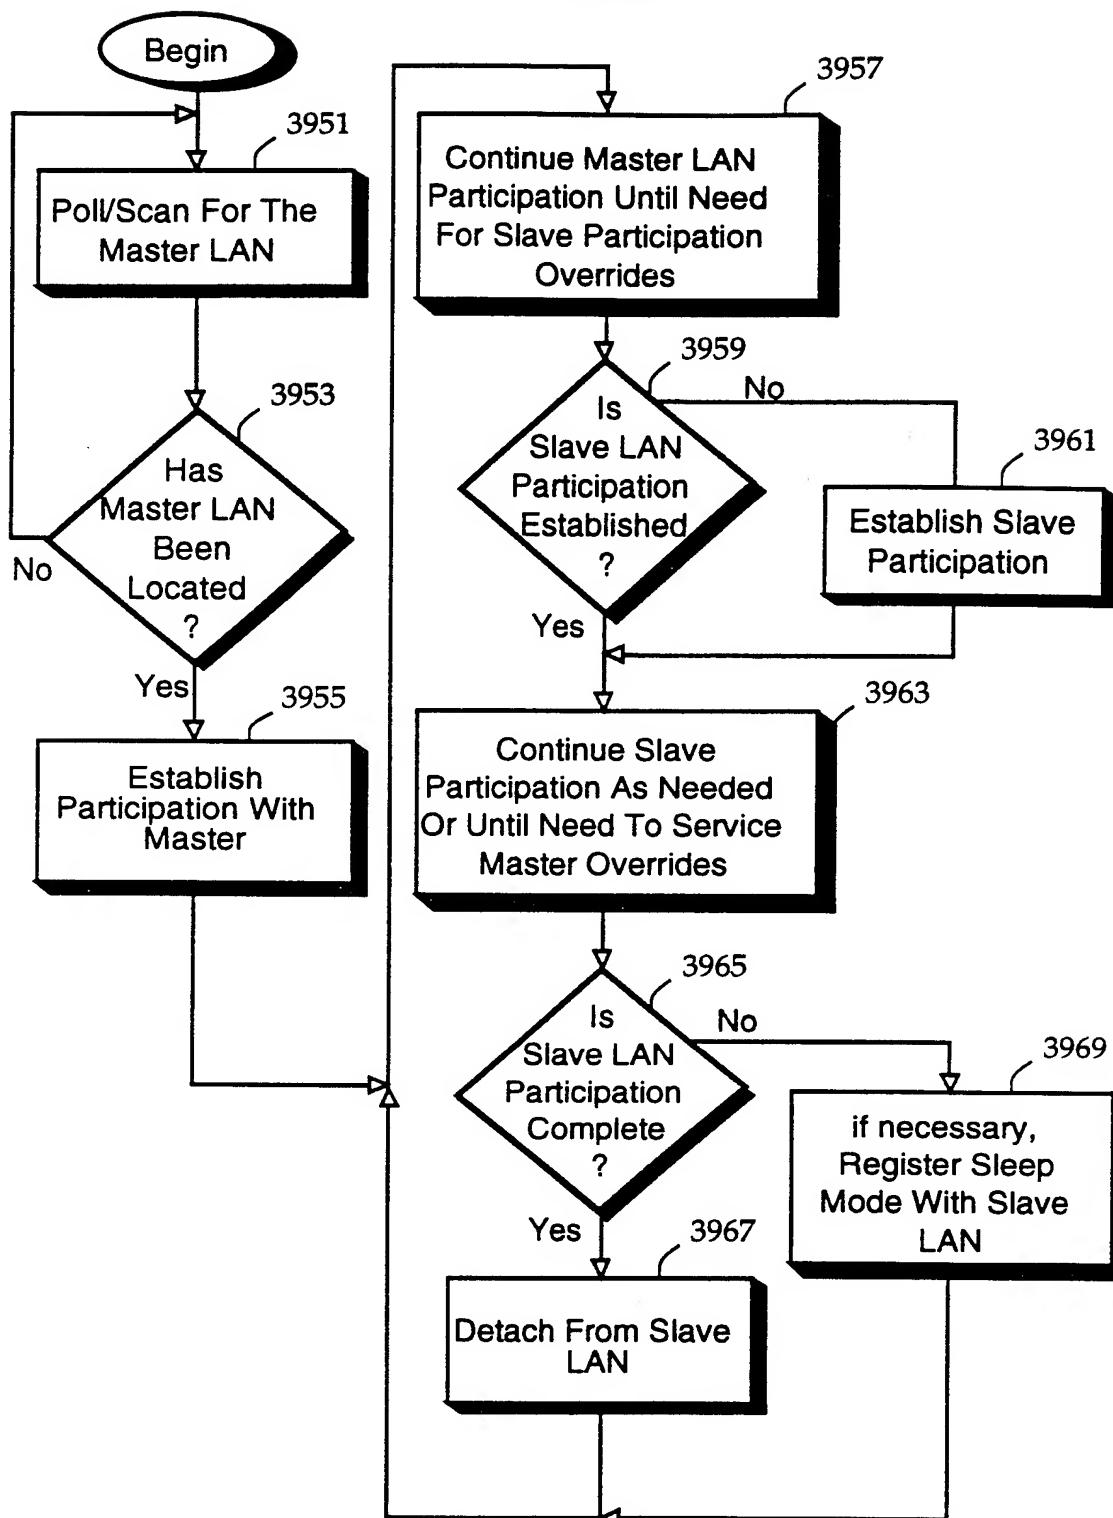


Figure 44

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US94/05037

A. CLASSIFICATION OF SUBJECT MATTER

IPC(5) :H04B 7/216; H04L 12/46

US CL :455/11.1, 56.1, 38.3, 343; 370/18, 84, 85.13, 95.1; 375/1

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 455/11.1, 33.1, 38.3, 53.1, 54.1, 56.1, 66, 343; 370/18, 84, 85.13, 95.1, 95.3; 375/1; 379/58, 59

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

APS

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US, A, 4,539,706 (MEARS ET AL.) 03 September 1985, col.4, lines 40-47.	1, 6, 9, 12, 18, 24, 26
A	US, A, 4,606,044 (KUDO) 12 August 1986, col. 3, lines 39-63.	4, 8
A	US, A, 5,029,183 (TYMES) 02 July 1991, entire document.	1-2, 6, 11-12, 18, 21, 24, 26
A, P	US, A, 5,220,678 (FEEI) 15 June 1993, entire document.	2, 6, 11, 24
A, P	US, A, 5,241,542 (NATARAJAN ET AL.) 31 August 1993, entire document.	2, 6, 11, 24
A, P	US, A, 5,295,154 (MEIER ET AL.) 15 March 1994, entire document.	1, 6, 9, 12, 18, 24, 26

☐ Further documents are listed in the continuation of Box C. ☐ See patent family annex.

* Special categories of cited documents:	*T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
A document defining the general state of the art which is not considered to be part of particular relevance	*X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
E earlier document published on or after the international filing date	*Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	*G* document member of the same patent family
O document referring to an oral disclosure, use, exhibition or other means	
P document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

09 SEPTEMBER 1994

Date of mailing of the international search report

SEP 20 1994

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

HASSAN KIZOU

Telephone No. (703) 305-4744